The Configurable Processor Company

Next-Generation Audio Engine

Robert Kennedy
Senior Software Engineering Manager
Darin Petkov
Member of Technical Staff

What's Changing in Audio

- Moving to System on Chip (SoC)
- Expanding universe of audio standards
- Portable and multi-purpose devices (e.g. handsets) feature:
  - Audio
  - Multiple non-audio applications
- Increasing pressure to reduce power and area

Audio Platform Requirements

- Low power and area
- Low cycle consumption by codecs
  - Multiple codecs run simultaneously
  - Cycles left for effects, mixdown, non-audio applications
  - Achieved through MAC, load/store, ALU, Huffman, control performance, parallelism, etc.
  - Fewer cycles -> lower clock rate -> lower power
- Flexibility / programmability
- Multiple data types (16- and 24-bit signal data, sometimes even 32-bit)
- Applicable to the widest range of audio products

Today's Approaches

- **General-purpose embedded CPU**
  - Not optimized for high-quality real-time sound processing
- **DSPs**
  - General purpose DSPs use more silicon area than required for audio applications
  - Not a good match for control tasks
- **Hard-wired RTL**
  - Requires one block per audio standard (makes the chip huge)
  - No changes possible without redesigning chip
- **Tensilica's HiFi 1 Audio Engine**
  - Based on Xtensa V architecture
  - Runs AC-3, G.723, G.729AB, MP3, MPEG-2/4 AAC and WMA
  - Designed into:
    - Cell phones
    - Portable Audio Players
  - With new Xtensa LX technology we do better

Xtensa LX makes HiFi 2 possible

- **Xtensa: Configurable, Extensible, Synthesizable**
  - Extensions driven by analysis of audio codecs
- **Enabling Xtensa LX features**
  - FLIX (Flexible Length Instruction eXtensions)
    - Base has 16- and 24-bit instruction sizes
    - Custom instructions can use 24-bit and 32- or 64-bit instruction sizes
    - 32- and 64-bit sizes allow multiple independent operations per instruction
    - FLIX relaxes single-issue programming model of Xtensa V / HiFi 1
  - Functional clock gating reduces power
- **HiFi 2 is Xtensa LX with a particular audio-specific set of instruction extensions**
  - More custom instructions can be added
  - Extensions are first-class citizens
  - Imposes a minimum configuration requirement

Instruction Set Overview

- HiFi 2 adds more than 300 operations
  - Dual multiply with 56-bit accumulate
    - Each multiplier supports 24 x 24 bits and 32 x 16 bits
    - Both multipliers operate every cycle
  - Add / subtract and variable / immediate shifts
  - Huffman encode / decode and bit stream support
    - Streams interleave coded / uncoded items
  - Convert / round / truncate instructions
  - Two special audio register files with multiple data types
    - P: 8 x 48 bits (each holds two 24-bit values)
    - Q: 4 x 56 bits (accumulator values)
  - Two way SIMD arithmetic and boolean operations on 24-bit or 16-bit data

Close-up view: MAC modes supported

- Single and dual multiplication
- Fractional and integer arithmetic
- Operands:
  - 24x24 bits P x P (typical audio)
  - 16x16 bits P x P with intermediate saturation (AMR, G.7xx)
  - 32x16 bits Q x P (WMA at low bit rates)
- Accumulation: overwrite, add, subtract; with or without saturation
- Signed and unsigned:
  - signed x signed (typical)
  - signed x unsigned (multiple precision)

Example MAC Instruction Usage

```
ae_p24x2s a, b;      /* allocated in P registers */
ae_q56s   x;         /* allocated in Q registers */
...
/* fractional real part of complex multiply:
 * x = a.H * b.H - a.L * b.L */
x = AE_MULZASFP24S_HH_LL(a, b);
```

Example MAC Instruction Mnemonic

```
x = AE_ MUL Z A S F P24S_ HH_ LL(a, b);
```

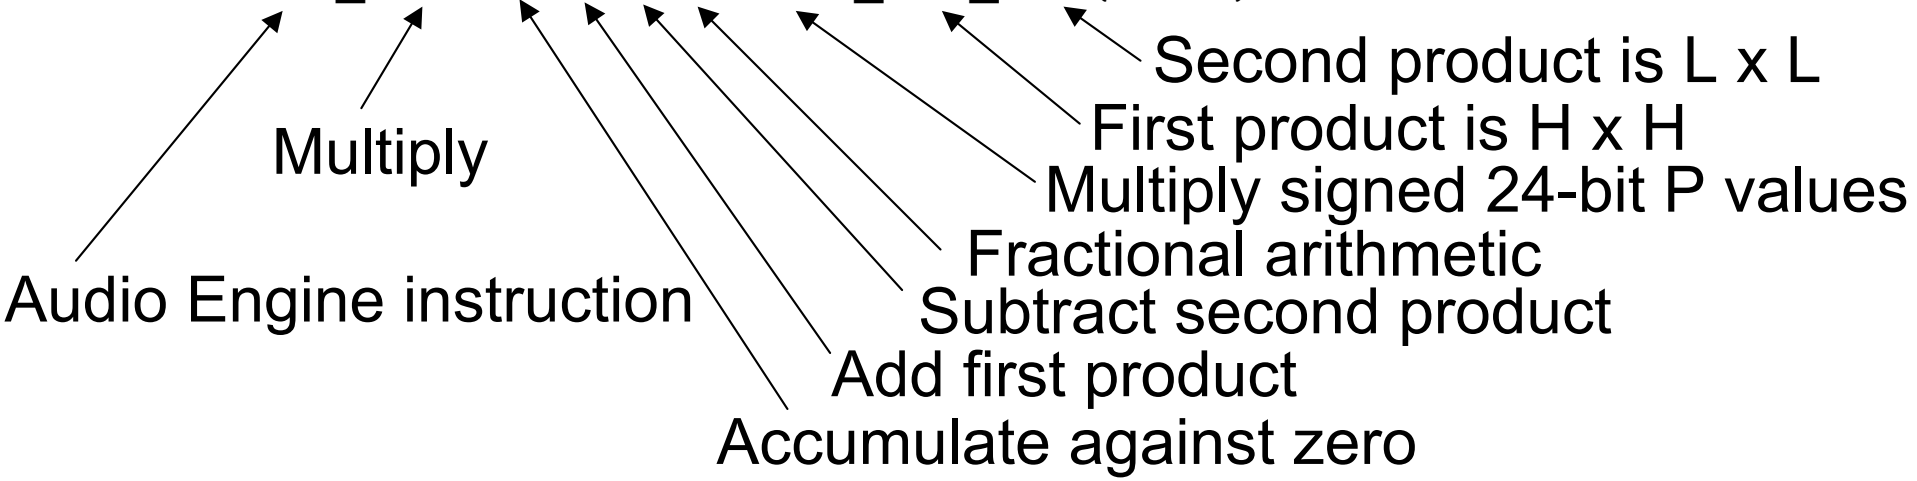

HiFi 2 Instruction Formats

FLIX: Flexible-Length Instruction eXtensions

Dual-Issue 64-bit FLIX or Single-Issue 24/16-bit Operations

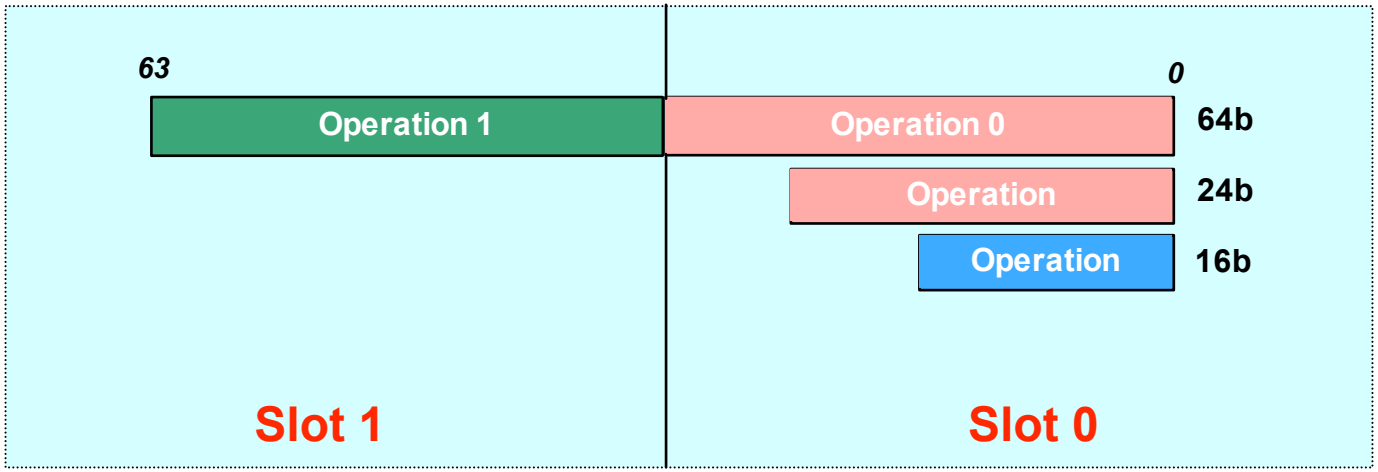

HiFi2 Audio Operations

Base / HiFi2 Audio Operations

Base LX Operations

Multiply and Audio ALU Operations

Load / store, Huffman Operations

HiFi2 Datapath and Registers

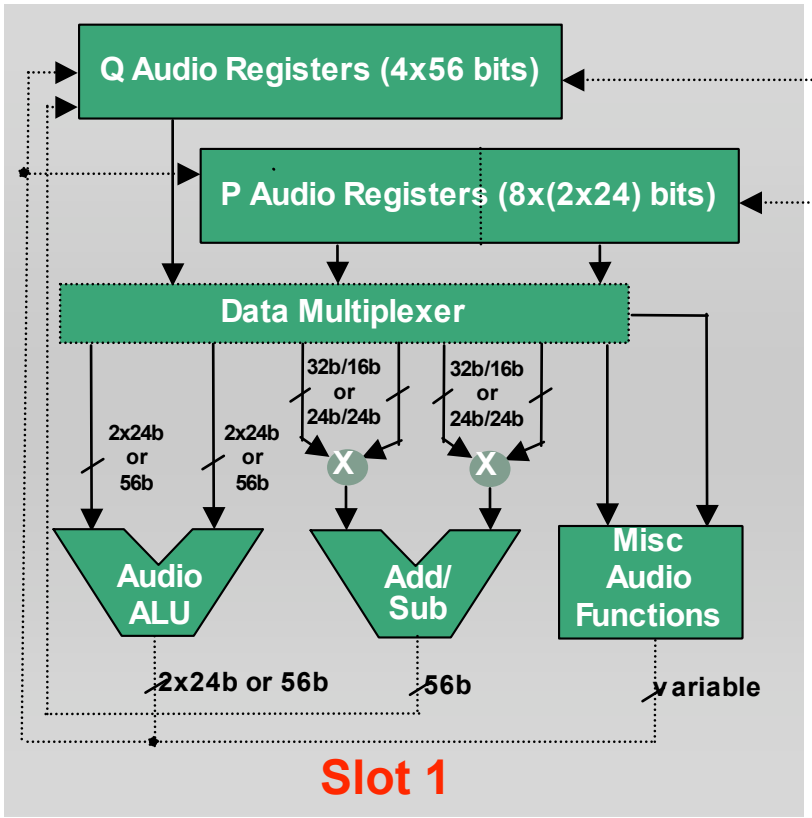

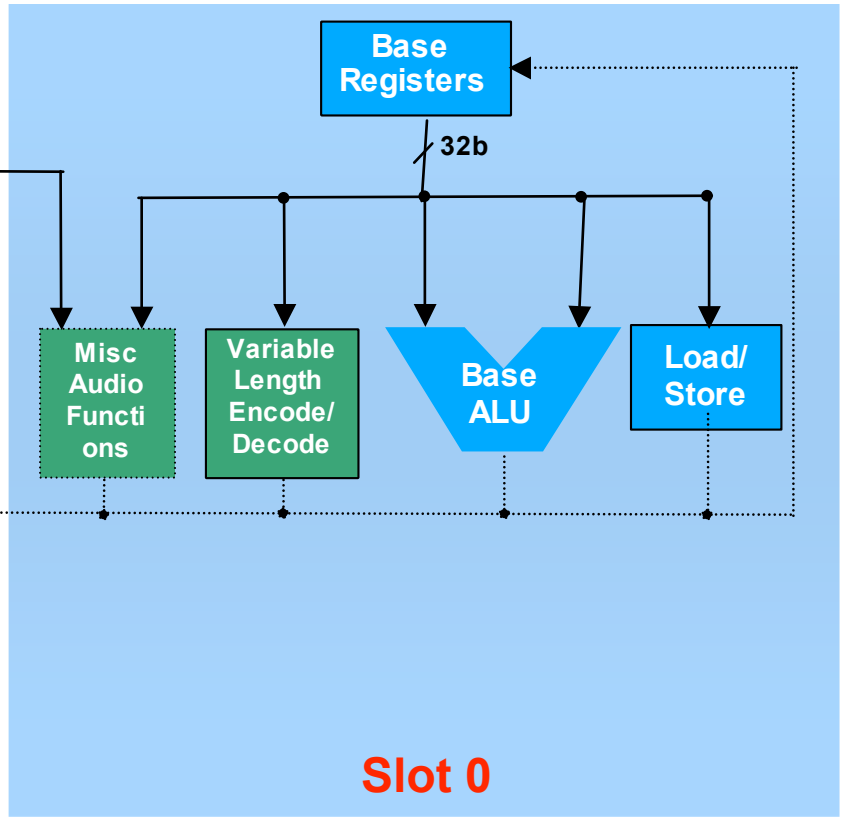

- HiFi 2 Audio Engine Hardware
- Base Xtensa LX Configuration

Design Alternatives Considered

- HiFi 2 MAC alternatives:
  - **24 x 24 bits (48-bit product, 56-bit accumulation)**
  - **32 x 16 bits (48-bit product, 56-bit accumulation)**
  - 32 x 32 bits (8 product bits discarded)
  - Single or **dual multiplier**
- Memory bandwidth:
  - **64**- vs. 128-bit bus requirement
  - **One** vs. two load/store units
  - **Bandwidth >2 GB/sec**

Implemented features shown in bold green

Configurations for Area, Speed, and Power Comparisons

- HiFi2 extensions
- 64-bit interface to memory
- 8k icache, 8k dcache, 2-way
- MUL32 option (~5-6k gates) present in one experiment

Example Configuration 1 Experiment: MAC Options and Hardware Cost

	Maximum clock rate (MHz)*	Gates*	Area* (mm^2)
Single 24x24-bit MAC	299	88,569	0.98
Dual 24x24-bit MAC	289	100,860	1.12
Dual MAC supporting 24x24 and 32x16	284	101,408	1.13
Dual MAC supporting 24x24, 32x16 and single 32x32	270	110,012	1.22

* Based on TSMC 0.13µ LV, Artisan library, includes MUL32 option

Example Configuration 2 Experiment: Power Dissipation Estimates in Simulation

Implementation	Area (mm^2)	Leakage power (mW)*	Switching power (mW/MHz)*	Real-time MP3 decode power (mW @ 14 MHz)
0.13µ lv** synthesized to 200 MHz	0.94	0.4	0.09	1.6
0.13µ g*** synthesized to 50 MHz	0.85	0.3	0.07	1.3

MUL32 option not present

* Power measured running MP3 decode

** Artisan SAGE-X library

*** Artisan metro library

Development Cycle Summary

- Six weeks from concept to first customer delivery
- Development guided by:
  - Software and hardware optimization experiments
  - Customer input
- Automatic processor generation provides:
  - Processor core RTL
  - Complete software tools
    - C/C++ compiler
    - Debugger
    - Linker
    - Simulator
    - Assembler
    - Profiler
    - RTOS Hardware Abstraction Layer

HiFi2 Design Approach Advantages

- Rapid turnaround (minutes) of tools generation and hardware estimation allows rapid assessment of
  - Software performance benefit
  - Hardware impact
- Full RTL generation in under an hour
- Software porting and optimization can (and should!) proceed concurrently with instruction set definition
- Optimized code uses
  - HiFi2-specific data types, register-allocated automatically by the compiler
  - HiFi2-specific instructions, generated by the compiler via instruction intrinsics
  - No assembly language (sure you can, but why?)

Selected Codec Preliminary Specs

Codec	Worst Case Required MHz
HiFi 2 MP3 Decoder	15-17
HiFi 1 MP3 Decoder	18
HiFi 2 MP3 Encoder	38-40
HiFi 1 MP3 Encoder	65
HiFi 2 AAC-LC Decoder	13-14
HiFi 1 AAC-LC Decoder	26
HiFi 2 AAC-LC Encoder	40-44
HiFi 1 AAC-LC Encoder	85
HiFi 2 WMA Decoder	18-21
HiFi 1 WMA Decoder	30

Summary

- Realistic configurations approaching 300 MHz, below 100k gates, below 1.5 mW for MP3 decode
- Excellent performance on broad set of audio applications, including future codecs
- Rich audio instruction set with complete, extension-aware software tools support
- Processor remains configurable to take on additional tasks
- Power, performance, and broad codec support make HiFi2 appropriate for a wide range of consumer and automotive products.