

Developing A/V Software for Consumer Media Products

Optimized DSP Software • Independent DSP Analysis




Developing A/V Software for Consumer Media Products (Workshop 266)

Berkeley Design Technology, Inc.
2107 Dwight Way, Second Floor
Berkeley, California 94704
USA
+1 (510) 665-1600


info@BDTI.com
<http://www.BDTI.com>

© 2003 Berkeley Design Technology, Inc.

Outline



Workshop Outline

The consumer media device 

- The big picture

Developing A/V software

- Software subsystems
- What's special about codec software?
- Numeric considerations
- Optimization techniques

Testing

Trends and conclusions

© 2003 Berkeley Design Technology, Inc. 2

© 2003 Berkeley Design Technology, Inc.



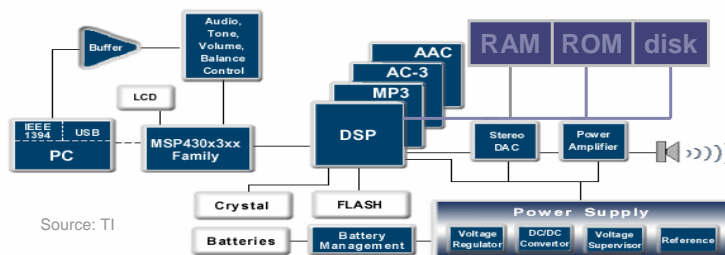
Big Picture

Consumer media device is complex system; several software and hardware subsystems to integrate, e.g.,

- Software: player, codec(s), RTOS
- Hardware: GPP/DSP, DMA, I/O, memory



© 2001 Iomega Corporation



Source: TI

© 2003 Berkeley Design Technology, Inc.

3



Key Big Picture Considerations:

Processor selection

- Numeric support: floating-point or fixed-point
- Development tools, e.g., IDE, HLL compiler, assembler
- Availability of OS, off-the-shelf SW, programmers, etc.

Reference implementations

- Hardware reference design
- Software components (RTOS, codec, player)

Design, development, and testing strategies


- Optimization
- Development board for early testing of software
- How testing will be performed, what resources will be utilized, and to what end: i.e., what's good enough

Integration

- Hardware + software
- Real-time testing

© 2003 Berkeley Design Technology, Inc.


4

Outline 

Workshop Outline

The consumer media device

- The big picture


Developing A/V software 

- Software subsystems
- What's special about codec software?
- Numeric considerations
- Optimization techniques

Testing

Trends and conclusions

© 2003 Berkeley Design Technology, Inc. 5

Developing A/V Software: Software Subsystems 


Software Subsystems

Primary software subsystems include:

- Player: GUI (play, stop, rewind), host helper func's**
- Codec(s): MPEG-2, WMA, RV9**
- I/O: DAC, USB**
- Real-Time Operating System: VxWorks, WinCE, Palm**

© 2003 Berkeley Design Technology, Inc. 6

Developing A/V Software: Software Subsystems



Key Software Considerations

Player

- Port to target OS/hardware platform (Just compile? Unlikely)

Codec

- Starting point?
- Floating-point to fixed-point migration (if necessary)
 - Numeric considerations
- Optimize for speed, memory use, power, etc.

RTOS

- Add/remove features/device drivers

Software integration


- Player + codec(s) + RTOS

Testing

- Audio/video quality (test vectors)
- Real-time performance

© 2003 Berkeley Design Technology, Inc. 7

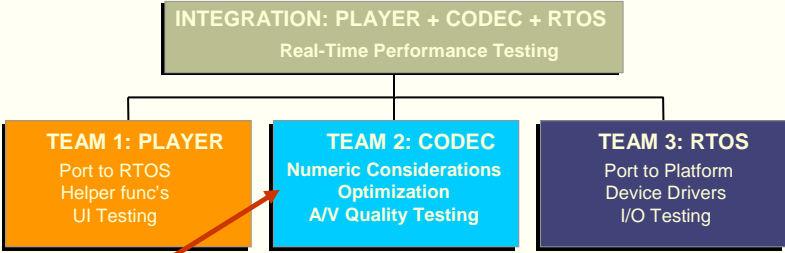
Developing A/V Software: Software Subsystems



Software Development

Common division of labor

- Separate teams for each subsystem
 - Teams work together to integrate and test




```
graph TD; I[INTEGRATION: PLAYER + CODEC + RTOS  
Real-Time Performance Testing] --- T1[TEAM 1: PLAYER  
Port to RTOS  
Helper func's  
UI Testing]; I --- T2[TEAM 2: CODEC  
Numeric Considerations  
Optimization  
A/V Quality Testing]; I --- T3[TEAM 3: RTOS  
Port to Platform  
Device Drivers  
I/O Testing];
```

Hot spot

- Codecs can pose significant development challenge

© 2003 Berkeley Design Technology, Inc. 8

Developing A/V Software: What's Special about Codec Software?



Codec Software Development


Not like other kinds of software development:

- Extreme computational demands
- Algorithm attributes
- Data access attributes
- Memory bandwidth requirements
- Testing and validation requirements
- Resource constraints
- Numeric fidelity requirements
- Standards
- Real-time requirements
- Reliability
- Specialized and complex processor architectures

→ Optimization is essential! ←

© 2003 Berkeley Design Technology, Inc. 9

Developing A/V Software: Numeric Considerations



Numeric Considerations

Many important and interesting topics, e.g.,

- **Float-to-fixed migration**
- **Numeric fidelity**
- **Data types**
- Error propagation/analysis
- **Precision and dynamic range**
- Block floating-point
- Floating-point emulation
- Signal scaling
- Rounding modes

(Focus topics)

© 2003 Berkeley Design Technology, Inc. 10



Float-to-Fixed Migration

Codec reference code available in different flavors, e.g.,

- Floating-point unoptimized
- Fixed-point unoptimized/optimized

Hardware platform usually fixed-point

- Fixed-point is cheaper, potentially faster, more energy efficient
- Few embedded processors support efficient floating-point

Floating-point reference code + fixed-point platform =

- Float-to-fixed migration

Float-to-fixed challenges

- Slows time-to-market
- Numeric fidelity tricky to maintain



Numeric Fidelity


Definition

- Numeric fidelity = accuracy of numbers

Why do we care?

- Computer arithmetic is not error-free
- Errors usually lead to:
 - Noise
 - Overflow—sudden change from max to min value
 - In video, e.g., white to black (bright to dark)
 - In audio, e.g., pop and click

How do we maintain numeric fidelity?


Developing A/V Software: Numeric Considerations 

Maintaining Numeric Fidelity

Minimize error & error propagation

- Analyze where errors (could) occur, and:
 - Determine tolerable level of error, e.g., one unit in last place
 - Choose alternative algorithm topology
 - To alter dynamic range requirements & error propagation
 - Choose appropriate data types and sizes
 - With sufficient dynamic range & precision for your signal
 - Scale signals
 - To prevent overflow or maintain precision
 - Select rounding modes
 - To minimize error propagation

© 2003 Berkeley Design Technology, Inc. 13

Developing A/V Software: Numeric Considerations 

Data Type Choice

Understand the attributes & implications of data types

	Fixed point	Floating point: IEEE-754 SP
Precision	16 bit: 1 part in 64K 24 bit: 1 part in 16M 32 bit: 1 part in 4G	24-bit mantissa: 1 part in 16M
Dynamic range	16 bit: 96 dB 24 bit: 144 dB 32 bit: 192 dB	8 bit exp: 1500 dB
Ease of use	Tricky	Easy
Processor cost	Cheap	Expensive

© 2003 Berkeley Design Technology, Inc. 14

BDTi

Data Types: Fixed-Point

16-bit fixed-point

Integer Fractional

Mixed

SNR (dB) Precision decreases as signal level decreases

© 2003 Berkeley Design Technology, Inc. 15

BDTi

Data Types: Fractional Numbers

Signals are often represented as a fraction in a fixed-point number:
S.3 or Q.3

Sign Bit

1/2
1/4
1/8

4-bit fractional value

Frac Val=	Sign Bit: S	b ₂	b ₁	b ₀
0.875	0	1	1	1
0.75	0	1	1	0
0.625	0	1	0	1
0.5	0	1	0	0
0.375	0	0	1	1
0.25	0	0	1	0
0.125	0	0	0	1
0	0	0	0	0

© 2003 Berkeley Design Technology, Inc. 16

Developing A/V Software: Numeric Considerations BDTi

Data Types: Fractional Add, Multiply

Fractional add ($N + N \geq N\text{-bit}$):

- Behaves the same as integer add
- Generates no error, if no overflow

Multiply ($N \times N \text{ bits} \geq N \text{ bit}$)

Fractional Result:
upper word

Integer Result:
lower word

```

short X, Y, Z;
Z = X + Y;
    
```

```

short X, Y, Z;
Z = ((long) X * (long) Y) >> 15;
    
```

© 2003 Berkeley Design Technology, Inc. 17

Developing A/V Software: Numeric Considerations BDTi

Data Types: Floating-Point

1 bit 8 bits 23 bits

S **Exp** **Mantissa** 32-bit floating-point

Value = $S \times (1 + \text{Mantissa}) \times 2^{\text{Exponent}}$

SNR (dB)

144

0

-1500 0 Signal Power (dB)

Granular-Noise Region

High-Quality Region

Overload Region

Precision roughly constant over dynamic range

6 dB

© 2003 Berkeley Design Technology, Inc. 18



Numeric Support

Fixed-point DSPs typically support fractional and integer fixed-point data types in hardware

- Fractional multiplication includes shift

Fixed-point GPPs typically do not support fractional data types in hardware

- Shift must be explicit

GPPs usually 32-bit, DSPs usually 16-bit, with some 32 bit support

Fixed-point DSPs and GPPs can emulate floating-point, but usually at a high MIPS cost

Floating-point processors are more expensive, use more energy, and have slower clock rates

No fractional support in ISO/ANSI C (coming soon?)



Optimization

Possible performance metrics:

- **Execution speed**
 - Processor-independent/dependent optimizations
 - High level language (HLL) optimizations
 - Hand code assembly for best performance
 - Memory access optimizations
 - Avoid cache, or L1, "thrashing"
- Memory usage (code size and data size)
 - May conflict with optimizations for speed
- Power consumption
 - Minimize off-chip memory accesses

→ Profile → Analyze → Optimize



Profiling Goal: ID S-rate Operations

80/20 Rule:

- 20% of software responsible for 80% of execution time, and vice-versa

Functions can be classified based on invocation rate:

Class	Invocation rate
I-rate (Initialization rate)	≤ 1 time per second
K-rate (Control rate—parameter updates, etc.)	~10 to 1,000 times per second
S-rate (Sample rate)	~10,000 to 100,000 times per second

At each level:

- Rate differs by 2 to 3 orders of magnitude vs. adjacent levels
- Execution cost increases by 2 to 3 orders of magnitude

I-rate and K-rate efficiency not too important, but S-rate efficiency very important



S-rate Operation: FIR Filter Kernel

C implementation of FIR kernel

$$y[n] = \sum_{k=0}^{T-1} x[n-k]h[k]$$

```
N=40;
T=16;

for (n=0; n<N; n++) {
    for (k=0, SUM=0; k<T; k++) {
        SUM += x[n-k] * h[k];
    }
    y[n] = SUM;
}
```

Developing A/V Software: Optimizations BDTi

Analysis: Compiled ARM7 FIR Filter

<pre> L1.20 B L1.80 MOV a4,#0 MOV lr,#0 L1.32 B L1.60 SUB v1,ip,a4 MOV v1,v1,LSL #1 LDRSH v1,[a1,v1] MOV v2,a4,LSL #1 LDRSH v2,[a2,v2] ADD a4,a4,#1 MLA lr,v2,v1,lr L1.60 CMP a4,v3 BLT L1.32 MOV a4,ip,LSL #1 STRH lr,[a3,a4] ADD ip,ip,#1 L1.80 CMP ip,v4 BLT L1.20 </pre>		<div style="border: 1px solid black; background-color: yellow; padding: 5px; margin-bottom: 10px;"> ARMCC compiler known to be very good. </div> <pre> N=40; T=16; for (n=0; n<N; n++) { for (k=0,SUM=0; k<T; k++) { SUM += x[n-k] * h[k]; } y[n] = SUM; } </pre>
---	--	---

© 2003 Berkeley Design Technology, Inc. 23

Developing A/V Software: Optimizations BDTi

Analysis: Compiled ARM7 FIR Filter

<pre> L1.20 B L1.80 MOV a4,#0 MOV lr,#0 L1.32 B L1.60 SUB v1,ip,a4 MOV v1,v1,LSL #1 LDRSH v1,[a1,v1] MOV v2,a4,LSL #1 LDRSH v2,[a2,v2] ADD a4,a4,#1 MLA lr,v2,v1,lr L1.60 CMP a4,v3 BLT L1.32 MOV a4,ip,LSL #1 STRH lr,[a3,a4] ADD ip,ip,#1 L1.80 CMP ip,v4 BLT L1.20 </pre>	<div style="border: 1px solid black; background-color: yellow; padding: 5px; margin-bottom: 10px;"> 18 instructions in kernel </div> <div style="border: 1px solid black; background-color: yellow; padding: 5px; margin-bottom: 10px;"> 4 branch instructions </div> <div style="border: 1px solid black; background-color: yellow; padding: 5px; margin-bottom: 10px;"> 2 instruction load sequence </div> <div style="border: 1px solid black; background-color: yellow; padding: 5px;"> Single instruction equivalent: LDRSH v2, [a2, a4, LSL #1] </div>
---	---

© 2003 Berkeley Design Technology, Inc. 24

Developing A/V Software: Optimizations BDTi

Analysis:

Algorithmic Transformation: # 1

Reorder coefficients: single index for x & h

<pre> L1.32 SUB v1, ip, a4 MOV v1, v1, LSL #1 LDRSH v1, [a1, v1] MOV v2, a4, LSL #1 LDRSH v2, [a2, v2] ADD a4, a4, #1 MLA lr, v2, v1, lr L1.60 CMP a4, v3 BLT L1.32 </pre>	<pre> L1.32 SUB v1, ip, a4 MOV v1, v1, LSL #1 MOV v2, a4, LSL #1 LDRSH v1, [a1, v2] LDRSH v2, [a2, v2] ADD a4, a4, #1 MLA lr, v2, v1, lr L1.60 CMP a4, v3 BLT L1.32 </pre>
---	---

© 2003 Berkeley Design Technology, Inc. 25

Developing A/V Software: Optimizations BDTi

Analysis:

Algorithmic Transformation: # 2

Count down rather than up, branch if index ≥ 0

<pre> L1.32 SUB v1, ip, a4 MOV v1, v1, LSL #1 MOV v2, a4, LSL #1 LDRSH v1, [a1, v2] LDRSH v2, [a2, v2] ADD a4, a4, #1 MLA lr, v2, v1, lr L1.60 CMP a4, v3 BLT L1.32 </pre>	<pre> L1.32 SUB v1, ip, a4 MOV v1, v1, LSL #1 MOV v2, a4, LSL #1 LDRSH v1, [a1, v2] LDRSH v2, [a2, v2] SUBS a4, a4, #1 MLA lr, v2, v1, lr L1.60 CMP a4, v3 BPL L1.32 </pre>
---	--

© 2003 Berkeley Design Technology, Inc. 26

Developing A/V Software: Optimizations

BDTi

Implementing Optimizations

Combining:

- Branch reduction
- Single index off-set load
- Re-order coefficients
- Count down loop

11 instructions in total, compared to 18 in compiled code: ~40% less

```
Loop:
add r7, r2, r4, lsl #2      ; r7 points to in(i)
mov r6, r5                  ; j = ntabs-1
mov r10, #0                 ; sum = 0
innerLoop:
  ldr r8, [r7, r6, lsl #2]  ; r8 = in(i+j)
  ldr r9, [r3, r6, lsl #2]  ; r9 = coef(j)
  mla r10, r8, r9, r10      ; sum += in(i+j)*coef(j)
  subs r6, r6, #1           ; j--
  bpl innerLoop             ; loop until j < 0
str r10, [r1, r4, lsl #2]   ; out(i) = sum
subs r4, r4, #1            ; i--
bpl loop                    ; loop until i < 0
```

© 2003 Berkeley Design Technology, Inc. 27


Developing A/V Software: Optimizations

BDTi

Memory Access Optimization

Video Processing

- Video frame much bigger than typical L1 memory (cache or SRAM)



- Simplified frame buffer and L1 memory

© 2003 Berkeley Design Technology, Inc. 28

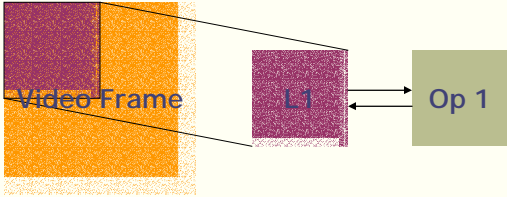
Developing A/V Software: Optimizations

BDTi

Memory Access Optimization

Default processing sequence

- Operation 1 on entire frame
- Operation 2 on entire frame
- Etc.



The diagram illustrates the default processing sequence. On the left, a large orange rectangle represents a 'Video Frame'. A smaller purple rectangle labeled 'L1' is shown to the right of the frame, with a double-headed arrow indicating data flow between them. To the right of the L1 cache is a green rectangle labeled 'Op 1', with a double-headed arrow indicating data flow between the L1 cache and the operation.

© 2003 Berkeley Design Technology, Inc. 29

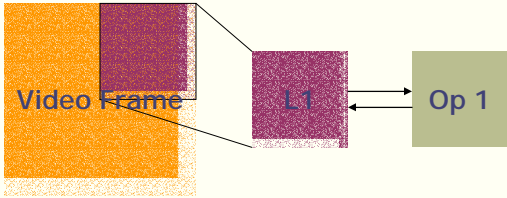
Developing A/V Software: Optimizations

BDTi

Memory Access Optimization

Load 2nd block of frame data

- Each cache line misses—very expensive—or
- DMA overhead for SRAM L1



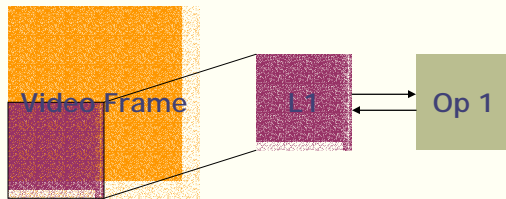
The diagram illustrates the loading of a second block of frame data. On the left, a large orange rectangle represents a 'Video Frame'. A smaller purple rectangle labeled 'L1' is shown to the right of the frame, with a double-headed arrow indicating data flow between them. To the right of the L1 cache is a green rectangle labeled 'Op 1', with a double-headed arrow indicating data flow between the L1 cache and the operation.

© 2003 Berkeley Design Technology, Inc. 30

Memory Access Optimization

Load 3rd block of frame data

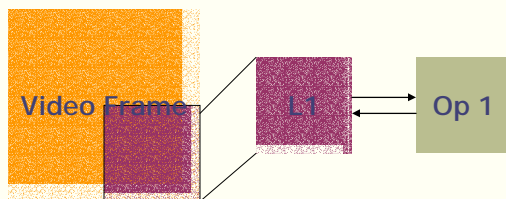
- Again, cache misses or DMA overhead
- Note: still Operation 1



Memory Access Optimization

Load Nth block of frame data

- Have moved entire frame in and out of L1
 - Slow and power-hungry

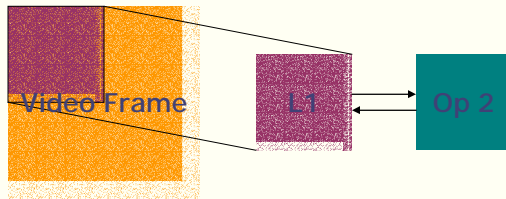




Memory Access Optimization

Repeat block load and process pattern for

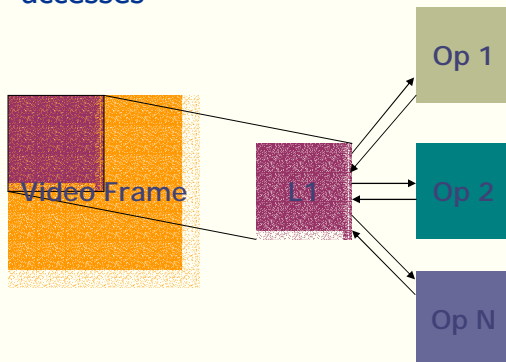
- Operation 2
- Ratio of memory access overhead to processing overhead very high




Memory Access Optimization

Optimization: process subset of frame

- Subset stays resident in L1, cutting external memory accesses



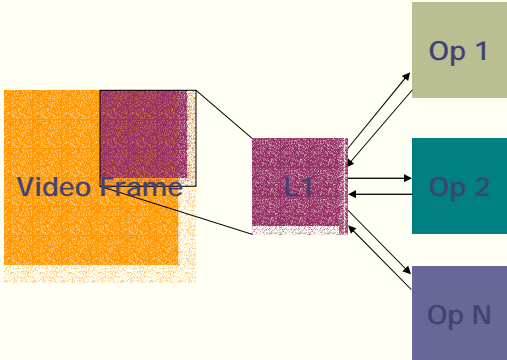
Developing A/V Software: Optimizations



Memory Access Optimization

Load next block of frame data, etc.


- Ratio of memory access overhead to processing overhead much lower



The diagram illustrates the flow of data from a video frame to processing operations. On the left, a large orange and purple rectangle represents a 'Video Frame'. A smaller purple rectangle labeled 'L1' (L1 cache) is shown to the right of the frame, with arrows indicating data being loaded from the frame into the cache. From the L1 cache, three arrows point to three separate colored boxes labeled 'Op 1' (green), 'Op 2' (teal), and 'Op N' (purple), representing different processing operations.

© 2003 Berkeley Design Technology, Inc. 35

Outline



Workshop Outline

The consumer media device

- The big picture

Developing A/V software

- Software subsystems
- What's special about codec software?
- Numeric considerations
- Optimization techniques

Testing ←

Trends and conclusions

© 2003 Berkeley Design Technology, Inc. 36

Testing 

Testing

Hardware/development platform

- Vast data I/O capability
 - Capture digital output for testing

Codec software

- Audio and video quality
 - Test vectors, reference codecs
- Operating modes
 - Sample rates, frame sizes, bit rates, etc.

System level (hardware + software)

- Real-time performance under worst cases
 - Data-dependent execution time
 - Dynamic processor features
 - Interrupts enabled

© 2003 Berkeley Design Technology, Inc. 37

Testing 

Hardware/Development Platform

A/V codecs consume and produce vast amounts of data:

- Compressed bit streams up to ~10 Mbps
- Uncompressed output streams up to 120 Mbps


Simulation model usually far too slow

- Real hardware is needed

Development board must have means to

- Supply large test vectors
- Capture potentially even larger output

© 2003 Berkeley Design Technology, Inc. 38

Testing 

Codec Software: A/V Quality


Difficult to measure quality in context of "lossy" compression algorithms

- Intentionally not bit-exact

Quality measured via reference codec + test vectors

- Supplied test vectors may not adequately stress fixed-point implementations
- May need to create tests that exercise full potential dynamic range of algorithm
 - Requires in-depth understanding of underlying algorithm

© 2003 Berkeley Design Technology, Inc. 39

Testing 

Codec Software: Operating Modes

A/V Codecs typically have several operating modes

- MPEG-1 Layer III audio has:
 - 14 bit rates
 - Three sampling rates
 - Four channel configurations

All valid combinations must be thoroughly tested

- Standard reference test vectors probably not sufficient
 - MPEG-1/2 Layer III audio has only one "compliance" test vector, so test vectors must be created

© 2003 Berkeley Design Technology, Inc. 40

Testing BDTi

System Level: Real-Time

Real-time performance is not optional
Processor is often underpowered

- Careful codec optimizations can pull underachievers up to real-time performance
 - But must verify worst-case performance

© 2003 Berkeley Design Technology, Inc. 41

Testing BDTi

System Level: Worst-Case Testing

Most demanding operating mode

- Highest bit rate, sample rate, most channels
- Interrupts enabled and active (UI and I/O)

Most demanding data


- Codecs have data dependent execution paths
- Processors often have data dependent operations, e.g., multiplication

Worst case real-time stress requires:

- Test vector which ensures worst case execution path
- Worst case inputs to data dependent operations

© 2003 Berkeley Design Technology, Inc. 42

Outline



Workshop Outline


The consumer media device

- The big picture

Developing A/V software


- Software subsystems
- What's special about codec software?
- Numeric considerations
- Optimization techniques

Testing

Trends and conclusions 

© 2003 Berkeley Design Technology, Inc. 43

Trends & Conclusions



Conclusions

Successful A/V codec software development:

- Demands knowledge of the application, algorithms, and processor, and mastery of a wide range of skills and tools
- Requires careful selection of numeric data types and close attention to numeric fidelity
- Typically requires aggressive optimization in order to meet tough real-time deadlines
- Requires a well-thought-out testing strategy!

© 2003 Berkeley Design Technology, Inc. 44



Trends

Processors are getting faster & compilers are getting better

- But newer A/V codecs are more demanding than previous generations

Optimized software libraries are more common

- Signal processing function level, e.g., FIR, IDCT
- Application level, e.g., A/V codec



Trends

Heterogeneous processors:

- Processor core + programmable logic
- Multi-processor SoCs

Heterogeneous processors may change how application workload is handled

- Proposal: off-load compute-intensive S-rate operations to custom logic or specialized DSP
- Reality: inter-processor communications and synchronization load can be deadly

Developing A/V Software for Consumer Media Products

Trends & Conclusions



Resources:

Processor Architecture/Software Optimization/OS:

Computers as Components—Principles of Embedded Computing System Design, Wayne Wolf

Numerics:

Digital Signal Processing, Alan V. Oppenheim & Ronald Schafer

Audio/Video/Speech Compression:

Digital Compression for Multimedia—Principles and Standards, Jerry D. Gibson, Toby Berger, Tom Lookabaugh, Dave Lindbergh, Richard L. Baker

© 2003 Berkeley Design Technology, Inc.

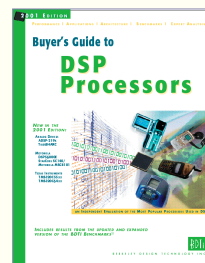
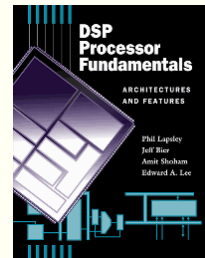
47

For More Information...

www.BDTI.com

Free Information

- White papers/presentation slides on
 - DSP software optimization
 - Streaming media implementation
 - Processor architectures and performance
 - Digital audio compression
- Article reprints on DSP-oriented processors and applications
 - *EE Times*
 - *IEEE Spectrum*
 - *IEEE Computer* and others
- *comp.dsp* FAQ



© 2003 Berkeley Design Technology, Inc.

2001 Edition

48

© 2003 Berkeley Design Technology, Inc.