



PORTABLE MULTIMEDIA

Powerful, Flexible Solutions Powered By CEVA's Mobile-Media

Author: **Eran Briman**, Sr. Director of Marketing

Contributors: **Erez Bar-Niv**, CTO DSP BU
Roni Sadeh, Multimedia Dept. Manager
Moshe Shahar, System Architecture Dept. Manager
Hagai Folkman, Senior Algorithm & Software Eng.

Table of Contents

Mobile Infotainment Sets Market Challenges	2
Wireless Multimedia Requires Flexibility	3
Mobile-Media: A Fully Programmable Complete Solution	4
Media Acceleration Technology in Depth	5
FST™- Replaces Conventional Software Implementations	6
Dedicated Multimedia DSP Capabilities Speed Building-Blocks	7
DMA Coprocessor Offloads Data Arrangement Completely	7
Conclusion	9

Mobile Infotainment Sets Market Challenges

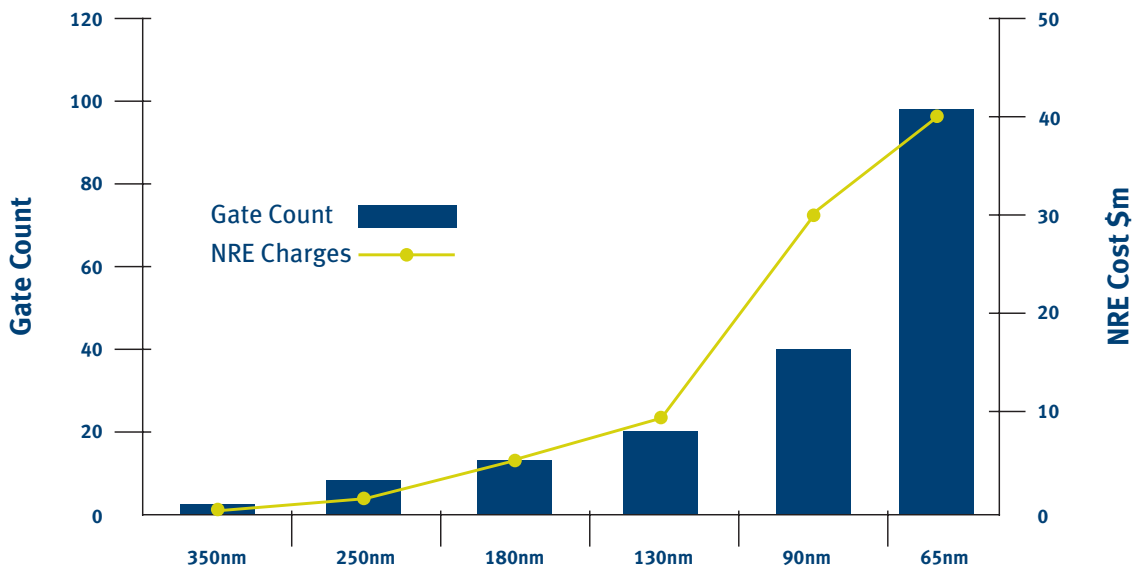
The convergence of mobility, connectivity and multimedia has created a market of enormous diversity and opportunity. From the network operators and content providers to the handset manufacturers and silicon vendors, companies strive to accommodate the subscriber's needs and wants – more multimedia features, higher data bandwidth, better screen resolutions, ubiquitous wireless connectivity – but not at the expense of small form factor and long battery life. In today's competitive market, this wealth of opportunities in mobile infotainment sets serious challenges for innovation and reusability.

Current market dynamics create a huge price pressure on every aspect of the supply chain. End users are reluctant to pay more than \$200 for fully loaded handsets, decreasing the margins for wireless carriers and handset manufacturers and forcing OEMs and ODMs to innovate with reusable platforms and technologies.


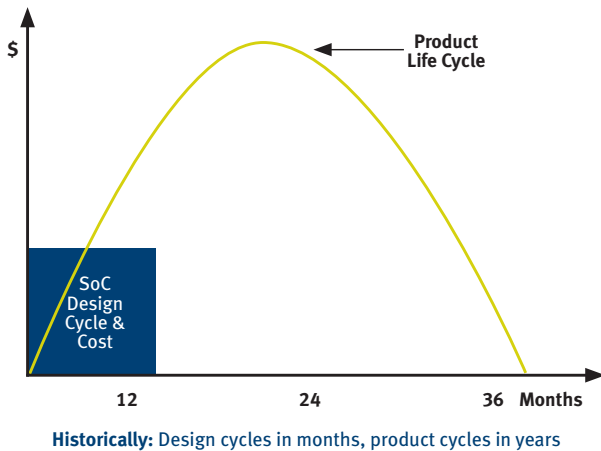
On the other hand, with increasing functionality in wireless devices, semiconductor content per phone continues to escalate, forcing semiconductor vendors to put in more silicon, higher development and integration costs and better process technologies. This poses enormous return-on-investment challenges on semiconductor vendors developing solutions for the mobile multimedia market. Semico Research forecasts that NRE costs for SoC designs in 65nm will approach \$40M per design. Increasing the challenge, end users expect to see a new mobile device in stores almost on a monthly basis, incorporating the latest advances in multimedia and connectivity. Product life cycles are becoming extremely short, generating less revenue per handset than ever before. For that reason, OEMs deeply segment their target markets while building reusable mobile platforms and re-spinning different derivatives with distinctive feature sets.

SoC Design Costs/Gate Count

Source: Semico Research, Feb 2005



The following diagrams demonstrate the extreme transformation in the governing market rules:



In summary, with today's markets constant pressure to produce cheaper and better devices, it simply costs too much and takes too long to effectively serve end users. The rapid product cycles and extended design cycles make the existing hardwired solutions hopelessly inadequate. This productivity gap calls for a new approach, one that would easily allow the reusability of the same platform for multiple devices, eliminating the need to re-spin silicon for every product and enabling future market trends to be easily accommodated.

A closer look at the multimedia applications in wireless devices, including mobileTV, video conferencing, capturing and showing video clips and still picture and listening to your favorite playlists, reveals an even more challenging business environment.

Wireless Multimedia Requires Flexibility

Wireless multimedia devices, including multimedia phones, smartphones and wireless PDAs, need to strike a thin and sometimes elusive balance between power consumption, feature set and cost. This is particularly true for today's wireless market, benefiting from 3G networks, high level of system integration and low bit-rate video codecs, but fiercely fighting tape-out costs and decreasing battery life.

With regards to power consumption, multimedia-enabled devices need to provide at least 2 hours of continuous media playback with massive media processing, memory accesses and high resolution screens. For that reason, platforms need to be efficient in processing media content, and would usually aim at the smallest geometry and lowest power process available. Moreover, handset manufacturers are ever more focused on small form factors, including large screens and keyboards. This is further driving the system integration trend, with more circuits being placed on the same silicon die, resulting in power reduction. However, smaller form factors also limit battery size and capacity.

The continuous struggle of cost versus functionality, with multimedia devices being mass market products, leads to a long list of functions with a very tight silicon budget. The ability to reuse the same hardware platform for all different functions in a given device, without having to spend more silicon real-estate, would be a win-win solution for this conflict.

When balancing the cost, power and feature set of a multimedia wireless devices, another significant factor needs to be considered. The coexistence of numerous media standards serving various multimedia applications (see table below) requires a platform that can be flexible enough to accommodate all of them. Furthermore, building a platform that can serve future standards that are not yet finalized is an even superior strategy.

Multimedia application	Content type	Standards used
Video playback & record, MMS	File-based	H.264, MPEG4, WM9, RealVideo
Videophone	Streaming	H.263, H.264
Streaming video, Mobile TV	Streaming, Interactive	H.264, MPEG4
Imaging	File-based	JPEG
Audio playback & record	File-based, streaming	MP3, WMA, AAC, aacPlus, BSAC, other

An approach that can accommodate this mix of standards and serve multiple functions on the same platform is to create a fully programmable solution. Such a software-based solution enables ultimate flexibility, supporting any given standard with no extra silicon costs and no additional risk. An already silicon-proven technology combined with the software environment can be easily reused for different devices across market segments and consumer audiences, from an entry-level feature phone to the most sophisticated wireless PDAs and smartphones.

Suppose this fully programmable solution would also meet or even exceed market requirements concerning performance, low power and cost targets.

Traditionally, there have always been issues with the concept of programmable multimedia platforms. Their processing power is usually not enough to support the latest video coding standards at high bit rates and such platforms are usually very power hungry, draining the battery too quickly, especially when high quality video is required. With CEVA's latest offering, a fully programmable solution can, for the first time, surpass these hurdles.

CEVA Inc., the leading licensor of DSP cores and communications solutions to the semiconductor industry, is now offering Mobile-Media, the first ever fully programmable solution for the mobile entertainment market. This complete solution supports H.264 encoding and decoding at up to full D1 resolution (720x480) at 30fps, without any hardware acceleration, and at an extremely low power envelope. CEVA has

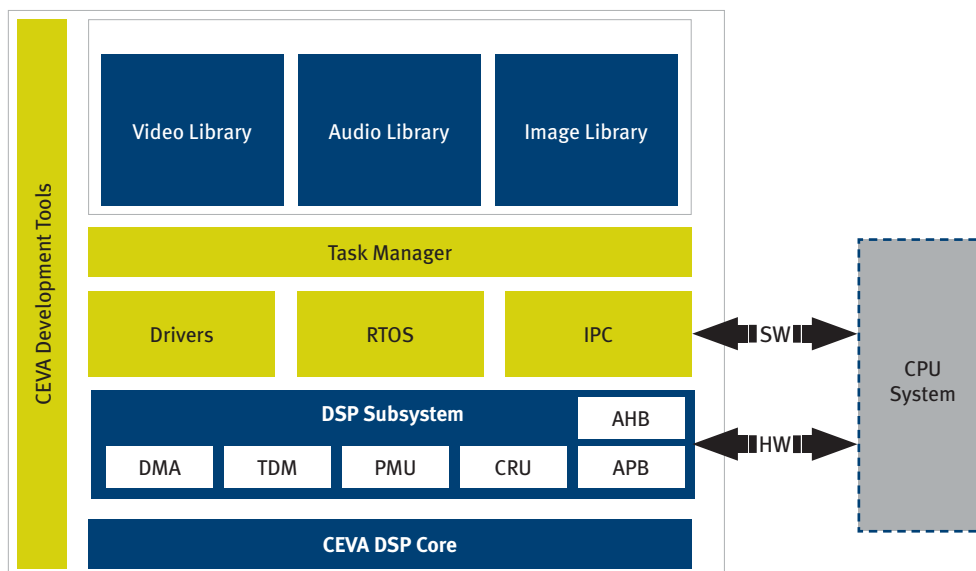
developed a patented SW-based media acceleration technology that significantly lowers computational power required for essentially any multimedia codec standard.

Mobile-Media - A Fully Programmable Complete Solution

Mobile-Media is a complete solution for mobile multimedia devices, including a DSP core, the surrounding subsystem, a software development environment and tools, and a full suite of optimized application software modules. The block diagram below depicts this complete platform.

Mobile-Media offers a single platform that serves all multimedia functions, including Video (different codecs at various resolutions and frame rates), Audio (different codecs), Imaging (JPEG at various resolutions and image enhancements) and Voice. By using pre-integrated AMBA bus bridges and standardized APIs, Mobile-Media allows licensees to simply integrate the platform into any given application processor platform, further reducing development effort and cost to a minimum.

Mobile-Media also drives down die size and associated cost, as the same piece of silicon is utilized for different purposes, without using any standard-specific hardware accelerators that can hardly be reused for other standards. Instead of integrating separate engines for audio, video, imaging and voice, a single platform such as Mobile-Media can serve all functions, while easing hardware and software integration complexities.



Mobile-Media includes two products: Mobile-Media1000 and Mobile-Media2000. Mobile-Media1000 is targeted at low-end to mid-range multimedia phones, supporting video streams of up to CIF resolution, 30 frames per second (fps) and is based on the CEVA-Teak DSP core. Mobile-Media2000 is targeted at the high-end multimedia phones, smartphones and Portable Media Players (PMPs) with a screen resolution of up to D1 30fps, and is based on the CEVA-X1620 DSP core. The table below summarizes the key features of these products. Note that some of the codecs and features in this table are optional and could be supplemented to the basic Mobile-Media packages.

Component	Platform	
	Mobile-Media1000	Mobile-Media2000
DSP Core	CEVA-Teak	CEVA-X1620
System Platform	Xpert-Teak	CEVA-XS1200
Development Environment	SW framework including APIs, RTOS and CPU communication layer, SW & HW development tools	
Target Market	Low-end to mid-range multimedia phones	High-end multimedia phones, and portable media players
Video	MPEG4, H.264, H.263	MPEG4, H.264, H.263 (Roadmap - WMV9, RV)
	Up to CIF, 30fps, encoding & decoding	Up to D1, 30 fps, encoding & decoding
	Simple profile (MPEG4)/Baseline profile (H.264)/ Profile O (H.263)	
Audio	AAC, MP3, WMA, aacPlus	
Imaging	JPEG (up to 1.3M pixel/sec), image enhancement	JPEG (up to 12M pixel/sec), image enhancement
Voice	AMR, WB-AMR, G.7xx	

In order to enable a completely programmable solution that does not make any usage of hardwired accelerators, CEVA has developed a software acceleration technology. Using it, Mobile-Media can easily run H.264 decoding at D1 resolution 30 fps, without consuming more than 100MHz and with no hardwired acceleration blocks. This acceleration technology is discussed in more detail in the following section.

Comparing CEVA's Mobile-Media solution with other fully programmable engines, such as TI's C55x, ARM9E and the recently announced ARM NEON, showcases the clear performance gains, mainly attributed to this unique multimedia technology. The following table lists the MHz required for two different multimedia applications using these programmable platforms¹.

Such a substantial gain, compared to other programmable platforms, is certainly a claim that needs to be further explored.

The multimedia technology behind this significant advantage, is a mixture of several multimedia technologies. It's not only the instruction set of the processor that is by far better tuned for media operations; it is also an algorithmic acceleration method, as well as a smart DMA engine managing all data transfers in the background.

	Mobile-Media 1000	Mobile-Media 2000	TI C5510	ARM9E	ARM NEON
MPEG4 SP encoder (QCIF, 15fps)	14	6	33	90	N/A
H.264 BP decoder (VGA, 30fps)	N/A	92	N/A	N/A	350

2

Media Acceleration Technology in Depth

Mobile-Media is employing an all-in-software acceleration technology, targeting the basic building blocks of different multimedia codecs. By accelerating the software building blocks, this technology can be easily applied to different multimedia standards (e.g. MPEG4, H.264 and WMV9) at different resolutions. These building blocks include:

- DCT and IDCT
- FFT
- Motion estimation
- Motion compensation
- Huffman coding, run length coding
- De-mosaic
- De-blocking filters
- Other Pre/post processing functions.

There are three main interconnected elements, forming this powerful technology:

1. FST™ – an algorithmic technique that helps to bypass most of the “brute-force” software implementations of the most demanding basic blocks.
2. Dedicated multimedia instructions and mechanisms, built into the programmable DSP, which further reduce the required computational horsepower and accelerate the software implementations.
3. DMA coprocessor built specifically for media content transfers, offloading the DSP from any type of data transformations and location mutations.

1. Sources: www.ti.com, www.arm.com

2. MHz requirements vary with bit rate and stream type.

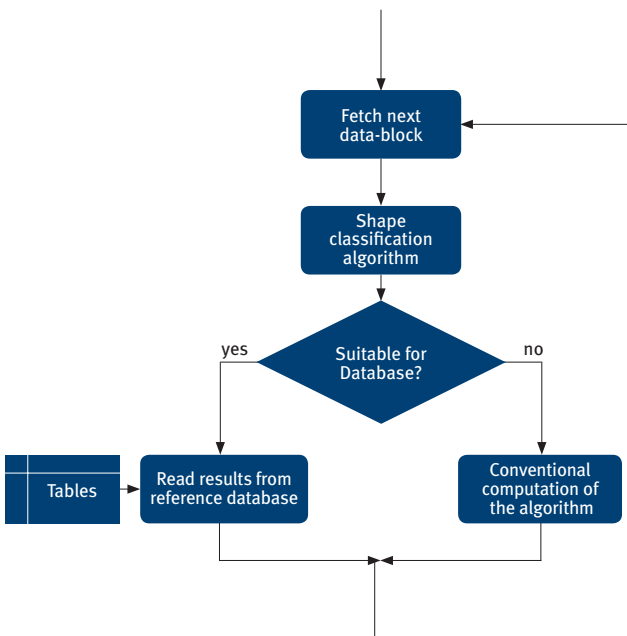
Each of these elements is briefly described below.

FST™- Replaces Conventional Software Implementations

FST is a unique algorithmic acceleration technique for the software implementation of multimedia building blocks, identifying and categorizing different shapes and indexes. By using this technique, “brute-force” software implementations of the same multimedia building blocks are accelerated by more than a factor of 10, resulting in a much faster codec implementation, and lower horsepower requirements.

FST uses a statistical model, recognizing specific shapes in the Macro-blocks being processed, and reading their pre-calculated results from prearranged set of tables that reside in the data memory. For every Macro-block (e.g. 8x8 block), an efficient and quick shape classification algorithm is first executed, in order to identify whether this Macro-block falls into a pattern that was recognized and has its results pre-stored in memory. This algorithm intensively calculates indices per Macro-block and also uses neighboring blocks’ information and attributes.

In the case where a shape is recognized, the conventional calculation is completely bypassed, and instead the results are read directly from the reference database in data memory. In the case where the shape classification algorithm does not identify a Macro-block with the same attributes, a conventional calculation of the building block is executed. The following diagram illustrates the flow of execution when using FST software acceleration.



FST allows, on average, to achieve a 90% hit ratio in the shape classification algorithm, while using relatively small data tables in memory – 32KB to 64KB. This means that, on average, 90% of Macro-blocks are recognized and go through the accelerated branch, fetching their pre-calculated results from the reference database. Whilst bypassing the conventional computation, the picture quality is not degraded in any way and the encoded and decoded streams remain totally compliant with the relevant standard (e.g. H.264).

The reference database should not necessarily reside in on-chip memory. Only small portions of these tables, the most frequently used, need to be embedded on-chip while the rest can be located in an external SDRAM.

FST algorithmic acceleration can be utilized in various multimedia building blocks. Following are 2 examples for the algorithmic acceleration achieved by utilizing FST:

MultiMedia Function	CEVA-Teak SW Implementation (in cycles)		FST Speed-up
	Conventional	FST-based	
8X8 DCT+ Quantization + Huffman Coding	2000	180	11x
Motion Estimation, 16X16 pixels macro-block, 128 pixel radius	15000	450	33x

FST brings innovation and state-of-the-art multimedia technology in 2 specific parts:

1. Building statistical tables that will require a very small amount of data memory, while providing a very high hit ratio.
2. Implementing the shape classification algorithm in software in a fast and efficient manner, avoiding the creation of new bottlenecks.

FST relies on several patents (some still pending), among them are:

- Data Compression Encoding System
- Method and apparatus for motion estimation in a sequence of digital images
- Method and apparatus for decompression of compressed data streams
- Method and apparatus for performing spatial-to frequency domain transform
- Enhanced Huffman Decoding + IDCT
- Enhanced Forward DCT + Huffman Encoding

Dedicated Multimedia DSP

Capabilities Speed Building-Blocks


FST can be applied to most multimedia building blocks and significantly accelerate them. However, considerable DSP multimedia capabilities are still required in order to complement FST's acceleration capabilities. These are used in conventional software implementations, such as in the following two scenarios:

1. Covering the rest of the 10% cases, where a known shape is not recognized and a conventional software implementation has to take place.
2. Multimedia building blocks that are not being accelerated by FST at all. The dedicated multimedia instructions are significantly accelerating those pure software implementations.

Below is a partial list of different instructions and mechanisms, embedded into the CEVA-X1620 DSP, specifically built to accelerate multimedia functions:

- Absolute differences – motion estimation, de-blocking filter
- Quad average – half/quarter pixel motion compensation
- Classify bytes – Non-linear filters, pre/post processing
- Byte addition/subtraction – DCT, motion reconstruct, quarter-pixel filters, symmetrical filters, motion estimation, de-blocking filter
- Inline data clipping (to byte or word dynamic ranges)– in-loop de-blocking filter
- Byte comparisons into accumulators – de-blocking filter
- Insert/extract – Huffman coding, RLC
- Minimum/maximum operations on bytes – de-blocking filter, non-linear filters
- Byte shifting with saturation – motion estimation, de-blocking filters
- Packing, unpacking, swapping – color conversion, pre/post processing

The following example illustrates the significance of these DSP multimedia instructions, in the software implementation of the Motion Compensation component of the H.264 decoding process.



The H.264 standard defines data blocks of 4x4 pixels, and a motion compensation process based on a quarter-pixel interpolation algorithm. In short, H.264 allows very fine motion vectors, shifting not only in pixels to any direction but also in half and quarter pixels. For that purpose, the motion compensation process of the decoder first has to generate 15 interpolated pixels for every 4 existing pixels, as shown in the diagram (G, H, M, N – 4 full pixels; b, h, j – 3 half pixels; a, c, d, e, f, g, i, k, n, p, q, r – 12 quarter pixels). Calculating these quarter pixels requires heavy averaging capabilities; some are simple averages, others are the result of a 6-tap 2D filter.

The CEVA-X1620 is equipped with the appropriate instructions and mechanisms to deal with these processing complexities. Among other features embedded in the DSP is an “average” instruction that enables the DSP to calculate up to 8 average bytes per cycle, and sophisticated byte additions and subtractions, treating different data sections as signed and unsigned values. The results of these arithmetic functions can then be saturated for different dynamic ranges, as specified in the H.264 standard.

DMA Coprocessor Offloads

Data Arrangement Completely

A DMA engine exists in all SoC designs for multimedia applications. The DMA's most important responsibility is to execute most data transfers, on- and off-chip, while accessing any available resource, including memories, I/Os, peripherals and bus bridges. By doing that, the DMA engine offloads parts of the data rearrangement tasks from the DSP, enabling it to focus mainly on media processing functions.

However, most DMAs are not capable of efficiently transferring video data, resulting in a large amount of overhead for the processor for managing the DMA engine, and even requiring

the processor to run some of the data transfers itself, calculate data address permutations and manipulate the data read from memory.

CEVA's Mobile-Media solution takes the DMA engine a step forward. The traditional DMA is replaced with a sophisticated DMA engine, which can be self-configured with minimal intervention required from the DSP. The basic contiguous source-destination buffers are replaced by dedicated channels for video frames, supporting 2D and 3D data transfers.

The DMA manager, supporting self-configuring channels, is mostly useful for dealing with non-sequential data, when different data sections from different locations must be gathered for processing repetitively. Instead of halting the DSP for reconfiguring the DMA every time new memory locations need to be programmed, the DMA coprocessor can reconfigure itself independently of the DSP, according to a task list the DSP prepares in advance.

One specific example for the enormous cycle savings gained by using the hardware DMA manager is the H.264 Motion Compensation function, discussed in the previous section. By nature, the Motion Compensation function requires the processing of many 4x4 Macro-blocks located anywhere around a given frame, depending on the motion vectors. Normally, the DSP would process a specific Macro-block while the DMA prepares the required data for the next block. Since this preparation requires gathering of different Macro-blocks, the DSP would be interrupted very often every time the DMA finished its transfer, in order to be reprogrammed for the next Macro-block transfer.

A	B	C	D	E	F
G	H	I	J	K	L
M	N	O	P	Q	R
S	T	U	V	W	X

In this diagram, the DSP is processing Macro-block U, while the DMA is preparing the data required for the next Macro-block to be processed – block V (note that the cubicles in this diagram represent Macro-blocks and not individual pixels).

For block V, according to its motion vectors, the DMA has to gather Macro-blocks E, F, K and L. Using conventional DMAs, the DSP would need to reprogram the DMA every time a new Macro-block has to be fetched by the DMA. A total of 4 interrupts would be required in the above example, with each interrupt adding between hundreds to thousands of cycles to the processing budget, just to reprogram the DMA engine.


Using the DMA manager, the DSP does not need to be interrupted at all. Instead, before starting the processing of block U, the DSP needs to copy the motion vectors of the next Macro-block, block V, to the HW DMA manager. The DMA manager can then analyze these motion vectors independently of the DSP and gather all the required Macro-blocks (E, F, K and L) one after the other. This unique capability significantly saves processing power, as the DMA offloads all data transfers from the DSP. Since the DMA obtains the exact Macro-blocks information, data traffic and L1 memory requirements are kept to a minimum and power consumption is reduced.

The DMA engine also includes dedicated 2D and 3D channels, specifically built to deal with the complex data transfers required for video frames. A very typical video data transfer scenario requires the collection of a specific Macro-block from several consecutive frames. Such pixel gathering from different frames is very intuitive for motion estimation functions, where a specific Macro-block needs to be compared to other Macro-blocks in previous or successive frames.

The dedicated 3D DMA channels allow the DMA to collect these Macro-blocks without any intervention from the DSP. By programming the DMA for different strides and different numbers of elements at three separate dimensions, the following transfer can be achieved totally independently of the DSP.

A conventional DMA engine would require complex programming and continuous involvement of the DSP in the data transfers, significantly increasing the MIPS count loaded on the DSP.

In summary, these three multimedia technology elements are the key enablers for a fully-programmable and flexible solution such as Mobile-Media. Dedicated multimedia instructions, such as the “sum of absolute difference” (SAD), commonly found in other DSP instruction sets, are far from adequate. Mobile-Media collects multimedia acceleration techniques in different areas of expertise, including DSP architectures, video algorithms and system optimizations, all running in harmony.



Video Frames Organization

DSP Memory

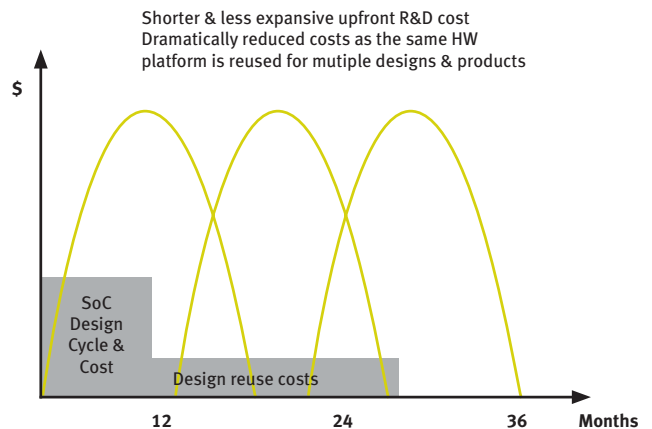
Conclusion

The ever increasing demands of multimedia processing capabilities from handset, PDA and other multimedia device manufacturers, and the escalating chip costs, are pushing the boundaries of silicon design. While hard-coded solutions were effective for one or two multimedia formats, the market is demanding that multiple formats and codecs be supported by the same design, to give a comprehensive and user-friendly experience to device owners. But as the audio, video and imaging formats became more numerous, the hard-coded paradigm was becoming cumbersome and ineffective.

The need for a fully programmable software solution is rapidly being understood by major handset manufacturers as the key to successful deployments of multimedia devices in the coming years. By reusing the same platform for multiple designs and devices, Mobile-Media bridges the gap between the escalating cost and design cycles, and the decreasing revenues and product life cycles, as shown below.

Powered by an innovative multimedia acceleration technology, this solution exceeds any criteria for mobile multimedia devices, including video quality, low power consumption and die size.

This solution can easily cope with all current multimedia standards and the myriad of evolving standards that have not yet been released to the market, attracting subscribers with stunning new multimedia-enabled devices, and driving up ARPU levels for wireless operators.



Mobile-Media enables design cycles in months targeting multiple product cycles



 www.ceva-dsp.com

PRINCIPAL OFFICES

USA

2033 Gateway Place, Suite 150, San Jose, CA 95110-1002,
Tel: +1 (408) 514 2900 Fax: +1 (408) 514 2995

Israel

2 Maskit Street, POBox 2068, Herzelia, 46120, Israel
Tel: +972 9 961 3700 Fax: +972 9 961 3800

Ireland

32-34 Harcourt Street, Dublin 2, Ireland
Tel: +353 1 402 5700 Fax: +353 1 402 5711

© 2005 CEVA, Inc. All Rights Reserved. All specifications are subject to change without notice. CEVA-Teak™, Xpert-Teak™, Xpert-Creator™, CEVA-X™, CEVA-X1620™, CEVA-XS™, CEVA-XS1200™ and Mobile-Media™ are trademarks or registered trademarks of CEVA, Inc. Other company and product names mentioned in this document may be the trademark or registered trademark of their respective manufacturers

Disclaimer: the information is provided "as is" without any express or implied warranty of any kind including warranties or merchantability, non-infringement of intellectual property, or fitness for any particular purpose. In no event shall CEVA, Inc. or its suppliers be liable for any damages whatsoever arising out of the use of or inability to use the materials. CEVA, Inc. and its suppliers further do not warrant the accuracy or completeness of the information, text, graphics or other items contained within these materials. CEVA, Inc. may make changes to these materials, or to the products described within.