TMS320C6713 Digital Signal Processor Optimized for High Performance Multichannel Audio Systems

Roshan Gummattira, Philip Baltz,
Nat Seshan

DSP Applications

ABSTRACT

The TMS320C6713's high performance CPU and rich peripheral set are tailored for multichannel audio applications such as broadcast and recording mixing, home and large venue audio decoders, and multi-zone audio distribution. The TMS320C6713 device is based on the high-performance advanced VelociTI™ very-long-instruction-word (VLIW) architecture developed by Texas Instruments (TI). The VelociTI architecture provides ample performance to decode a variety of existing digital audio formats and the flexibility to add future formats.

This paper will describe the following parts of the TMS32C6713 processor and their impact on high performance multichannel audio systems:

- The external peripheral architecture

- The C67x CPU architectural features and performance

- The real-time two-level cache architecture

- The multichannel audio serial ports (McASPs)

Contents

Trademarks are the property of their respective owners.

List of Figures

List of Tables

1   Introduction

High performance, multichannel audio applications are evolving at a rapid rate. In the consumer space, many standards have been defined. For example:

- Theater and home theater surround standards including: Dolby Pro Logic (II)®, Dolby Digital (EX)®, DTS(-ES)®, Sony Dynamic Digital Sound (SDDS®).

- Digital audio formats for portable and/or higher density (greater compression) playback: MPEG 2 Layer 3 (MP3), AAC, MPEG 4, Microsoft Windows Media, Meridian Lossless Packing(MLP)® (DVD-Audio), Rich Music Format (RMF)®.

In addition to consumer standards, many companies are developing their own high performance multichannel audio applications. Digital technology is being applied to large venues such as stadiums, auditoriums, and movie theaters to tune the listening experience to the room acoustics. Audio broadcast, production, and recording equipment implement effects generation as well as multichannel audio mixing, equalization, enhancement, and music.

1.1   System I/O

Figure 1 shows a block diagram of a digital surround receiver. Figure 2 generalizes that to many high performance multichannel audio systems. The TMS320C6713's peripheral set enables ease of connection to the major elements of these systems. The peripheral set includes:

- Two McASPs that provide simple cost effective, connectivity to multiple serial digital audio streams

- Two Inter-IC (I2C) buses for connection to serial ROMs or to control other system interface devices like user I/O

- A dedicated general-purpose input/output (GPIO) module to provide direct control lines to system components, eliminating much of the glue logic in many designs

- Glueless external memory interface (EMIF) capable of interfacing to SDRAM for bulk external storage of additional code or delay buffers. The EMIF also supports synchronous burst SRAM (SBSRAM), asynchronous memories, and peripherals with parallel interfaces.

- A host-port interface (HPI) for direct connection to a host processor

Figure 3 shows additional peripherals and the internal connection of the device. This includes:

- A highly efficient 16-channel enhanced direct memory access (EDMA) controller connects the peripherals to the internal and external memory. This controller can interleave transfers from different sources/destinations on a cycle-by-cycle basis, avoiding dead time of most DMAs when a higher priority transfer interrupts a lower priority one.

- Highly configurable PLL and clocking control logic to enable a variety of ratios of system and CPU clocks

- 256K bytes of internal memory to provide a large internal program and data store

- Two multichannel buffered serial ports (McBSPs) provide general connection to multiple serial standards including SPI

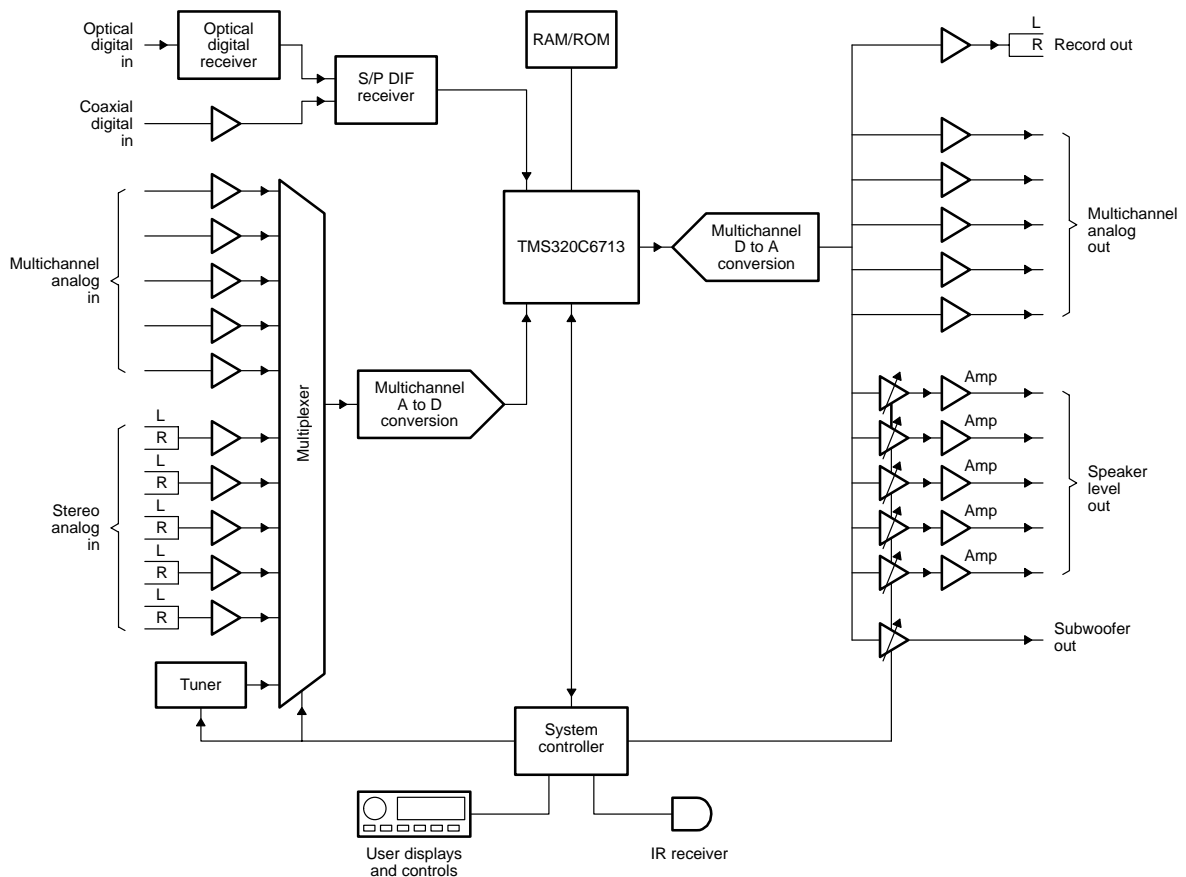- Two general-purpose timers to count system events or generate clock outputs

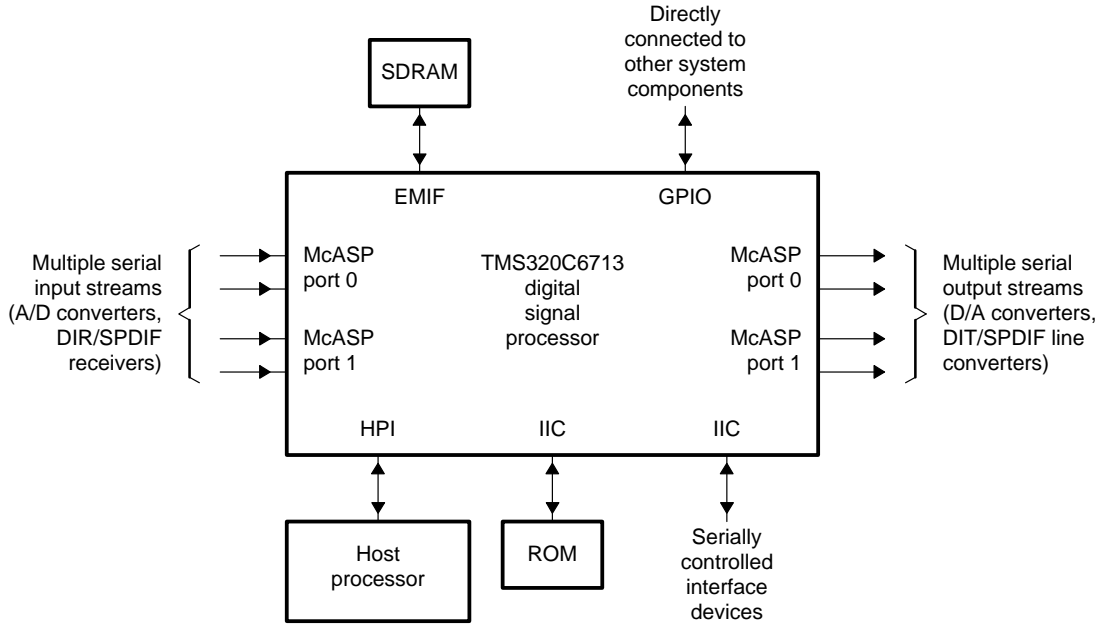**Figure 1. Digital Surround Receiver Block Diagram**

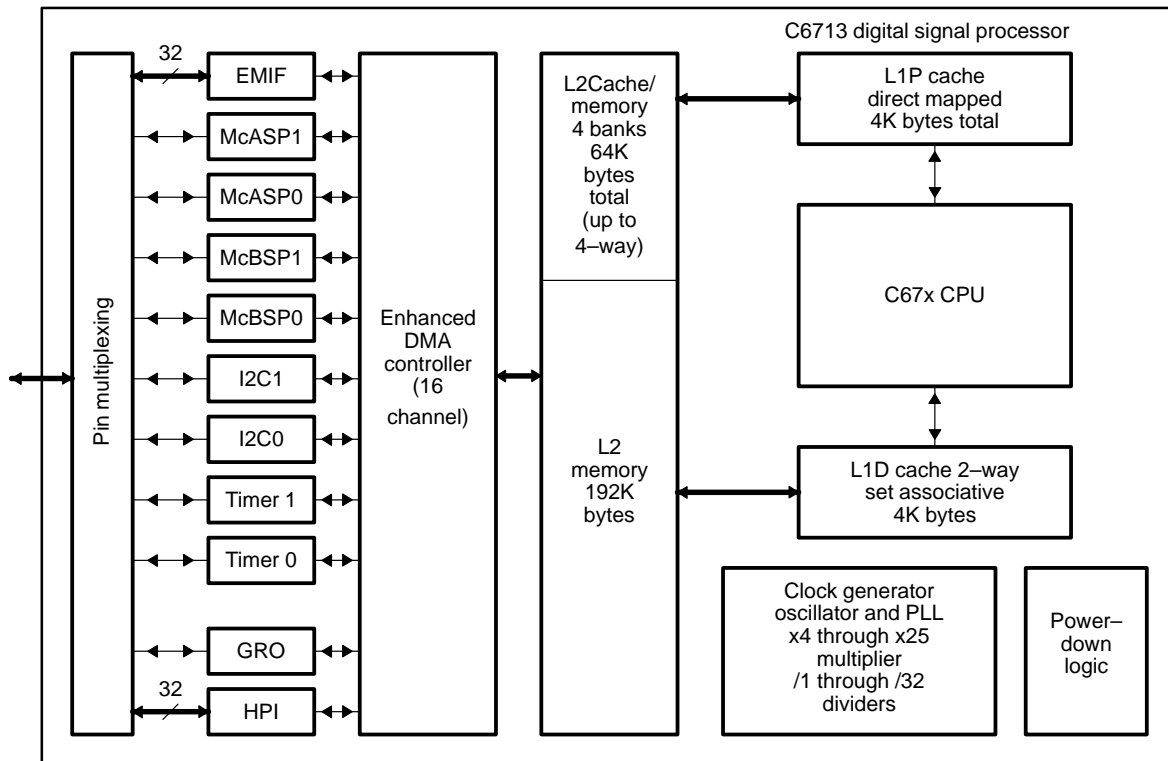**Figure 2. Generalized High Performance Multichannel Audio System**

Figure 3. TMS3206713 CPU and Peripheral Connectivity.

2   C67x CPU and Instruction Set

The TMS320C6713 floating-point digital signal processor uses the C67x VelociTI advanced very-long instruction words (VLIW) CPU. The CPU fetches (256 bits wide) to supply up to eight 32-bit instructions to the eight functional units during every clock cycle. The VelociTI VLIW architecture also features variable-length execute packets; these variable-length execute packets are a key memory-saving feature, distinguishing the C67x CPU from other VLIW architectures.

Operating at 225 MHz, the TMS320C6713 delivers up to 1350 million floating-point operations per second (MFLOPS), 1800 million instructions per second (MIPS), and with dual fixed-floating-point multipliers up to 450 million multiply-accumulate operations per second (MMACS).

2.1   Functional Units

The CPU features eight of functional units supported by 32 32-bit general purpose registers. This data path is divided into two symmetric sides consisting of 16 registers and 4 functional units each. Additionally, each side features a data bus connected to all the registers on the other side, by which the two sets of functional units can access data from the register files on the opposite side.

2.2   Fixed and Floating Point Instruction Set

The C67x CPU executes the C62x integer instruction set. In addition, the C67x CPU natively supports IEEE 32-bit single precision and 64-bit double precision floating point. In addition to C62x fixed-point instructions, six out of the eight functional units also execute floating-point instructions: two multipliers, two ALUs, and two auxiliary floating point units. The remaining two functional units support floating point by providing address generation for the 64-bit loads the C67x CPU adds to the C62x instruction set. This provides 128-bits of data bandwidth per cycle. This double-word load capability allows multiple operands to be loaded into the register file for 32-bit floating point instructions. Unlike other floating point architectures the C67x had independent control of the its two floating point multipliers and its two the floating point ALUs. This enables the CPU to operate on a broader mix of floating point algorithms rather than to be tied to the typical multiply-accumulate oriented functions.

2.3   Load/Store Architecture

Another key feature of the C67x CPU is the load/store architecture, where all instructions operate on registers (as opposed to directly on data in memory). Two sets of data-addressing units are responsible for all data transfers between the register files and the memory. The data address driven by the .D units allows data addresses generated from one register file to be used to load or store data to or from the other register file.

2.4   Benchmark Performance

Table 1 shows the TMSC32067x CPU floating-point benchmark performance of some algorithms commonly used in audio applications. The times for each benchmark are listed for a 225 MHz C6713 CPU.

Table 1. C6713 Benchmark Performance

Algorithm	Description	Parameter Values	Cycles	Time
Biquad filter (IIR filter direct form II)	nx input/output cycles	nx = 60 nx = 90	316 436	1.4 µs 1.9 µs
Real FIR filter	nh coefficients nr output samples	nh = 24 nr = 64 nh = 30, nr = 50	802 795	3.6 µs 3.5 µs
IIR filter	nr number of output samples	nr = 64	443	2.0 µs
IIR lattice filter	nr number of samples nk number of reflection coefficients	nk = 10, nr = 100	4125	18.3 µs
Dotproduct	nx number of values	nx = 512	281	1.2 µs

3   Two-Level Cache

3.1   Cache Overview

The TMS320C6713 device utilizes a highly efficient two-level real-time cache for internal program and data storage. The cache delivers high performance without the cost of large arrays of on-chip memory. The efficiency of the cache makes low cost, high-density external memory, such as SDRAM, as effective as on-chip memory.

The first level of the memory architecture has dedicated 4K Byte instruction and data caches, L1I and L1D respectively. The LII is direct-mapped where as the L1D provides 2-way associativity to handle multiple types of data. The second level (L2) consists of a total of 256K bytes of memory. 64K bytes of this can be configured in one of five ways:

- 64K 4-way associative cache
- 48K 3-way associative cache, 16K mapped RAM
- 32K 2-way associative cache, 32K mapped RAM
- 16K direct mapped associative cache, 48K mapped RAM
- 64K Mapped RAM

Dedicated L1 caches eliminate conflicts for the memory resources between the program and data busses. A unified L2 memory provides flexible memory allocation between program and data for accesses that do not reside in L1.

3.2   Cache Hides Off-Chip Latency

The external memories that interface to the TMS320C6713 may operate at a maximum of 100 MHz, while the device operates at a 225 MHz maximum frequency. All external memory devices have significant start-up latencies associated with them. For example, SDRAMs typically have a read latency of 2-4 bus cycles. The reduced frequency and additional latency of memories would normally significantly degrade processor performance. There is a significant reduction in latency for retrieving data from on-chip L2 memory than from an external memory. By having the intermediate L2 cache, this latency is hidden from the user. Using the fast L2 memories to cache the slower external memories reduces the latency of external accesses by a factor of five.

3.3 Unified L2 for Program and Data

By unifying the program and data in the L2 space, the L2 cache is more likely to hold the memory requested by the CPU. It enables the on-chip memory to contain more data than program when highly computational, looping code is being run to process large data streams. For long, serial code with few data accesses, the L2 may be more densely populated with program instructions. The unification allows you to allocate the appropriate amount of memory for both program and data and keeps the on-chip memory full of instructions and data that are the most likely to be requested by the CPU.

3.4 Real Time Features

An important concern in audio systems is that the device be able to perform in real time. There are several requirements for a system to ensure that real-time operation is possible. The operation of the device must be predictable, interrupts to the CPU must be handled without affecting the continued real-time operation of the device, and efficient I/O must be maintained.

3.4.1 Interrupt Handling

Interrupt handling is an important part of DSP operation. It is crucial that the DSP be able to receive and handle interrupts while maintaining real-time operation. In typical applications, interrupt frequency has not increased in proportion to the increase in device operation frequency. As processing speeds have increased, latency requirements have not.

The TMS320C6713 is capable of servicing interrupts with a latency of a fraction of a microsecond when the service routine is located in external memory. By configuring the L2 memory blocks as memory-mapped SRAM, or by using the L2 memory mapped space, it is possible to lock critical program and data sections into internal memory. This is ideal for situations such as interrupts and OS task switching. By locking routines that need to be performed in minimal time, the microsecond delay for interrupts is reduced to tens of nanoseconds.

3.4.2 Real Time I/O

Peripherals are a feature of most DSP systems that can take advantage of the memory-mapped L2 RAM. Typical processors require that peripheral data first be placed in external memory before it can be accessed by the CPU. The TMS320C6713 can maintain data buffers in on-chip memory, rather than in off-chip memory, providing a higher data throughput to peripherals. This increases performance when using on-chip McASPs, the HPI, or external peripherals. The EDMA can be used to transfer data directly into mapped L2 space while the CPU processes the data. This increases performance since the CPU is not stalled while fetching data from slow external memory or directly from the peripheral. Using this method for transferring data also minimizes EMIF activity, which is crucial as data rates or the number of peripherals increase.

3.5 Cache Summary

The efficiency of the cache architecture makes the device simple to use. The cache is inherently transparent to the user. Due to the level of associativity and the high cache hit rate, virtually no optimization must be done to achieve high performance. Reduced time for optimization leads to reduced development time, allowing functional systems to be up and running quickly. High performance can be immediately achieved with the cache architecture, while a Harvard architecture device with small internal memory requires much more time to achieve similar performance. This is because optimizing an application on a small Harvard architecture requires several iterations to tune the application to fit in the small, fixed internal memories.

4 McASP

4.1 McASP Overview

The McASP is a serial port optimized for the needs of multichannel audio applications. With two McASP peripherals, the TMS320C6713 device is capable of supporting two completely independent audio zones simultaneously.

Each McASP consists of a transmit and receive section. These sections can operate completely independently with different data formats, separate master clocks, bit clocks, and frame syncs or alternatively, the transmit and receive sections may be synchronized. Each McASP module also includes a pool of 16 shift registers that may be configured to operate as either transmit data, receive data, or general-purpose I/O (GPIO).

The transmit section of the McASP can transmit data in either a time-division-multiplexed (TDM) synchronous serial format or in a digital audio interface (DIT) format where the bit stream is encoded for S/PDIF, AES-3, IEC-60958, CP-430 transmission. The receive section of the McASP supports the TDM synchronous serial format.

Each McASP can support one transmit data format (either a TDM format or DIT format) and one receive format at a time. All transmit shift registers use the same format and all receive shift registers use the same format. However, the transmit and receive formats need not be the same.

The McASP has additional capability for flexible clock generation, and error detection/handling, as well as error management.

4.2 TDM Synchronous Transfer Mode

The McASP supports a multichannel, time-division-multiplexed (TDM) synchronous transfer mode for both transmit and receive. Within this transfer mode, a wide variety of serial data formats are supported, including formats compatible with devices using the Inter-Integrated Sound (IIS) protocol.

TDM synchronous transfer mode is typically used when communicating between integrated circuits such as between a DSP and one or more ADC, DAC, CODEC, or S/PDIF receiver devices. In multichannel applications, it is typical to find several devices operating synchronized with each other. For example, to provide six analog outputs, three stereo DAC devices would be driven with the same bit clock and frame sync, but each stereo DAC would use a different McASP serial data pin carrying stereo data (2 TDM time slots, left and right).

In the TDM synchronous transfer mode, the McASP continually transmits and receives data periodically (since audio ADCs and DACs operate at a fixed-data rate). The data is organized into frames.

In a typical audio system, one frame is transferred per sample period. To support multiple channels, the choices are to either include more time slots per frame (and therefore operate with a higher bit clock) or to keep the bit clock period constant and use additional data pins to transfer the same number of channels. For example, a particular six-channel DAC might require three McASP serial data pins; transferring two channels of data on each serial data pin during each sample period (frame). Another similar DAC may be designed to use only a single McASP serial data pin, but clocked three times faster and transferring six channels of data per sample period. The McASP is flexible enough to support either type of DAC but a transmitter cannot be configured to do both at the same time.

For multiprocessor applications, the McASP supports a large number of time slots per frame (between 2 and 32), and includes the ability to 'disable' transfers during specific time slots.

In addition, to support of S/PDIF, AES-3, IEC-60958, CP-430 receivers chips whose natural block (McASP frame) size is 384 samples; the McASP receiver supports a 384 time slot mode. The advantage to using the 384 time slot mode is that interrupts may be generated synchronous to the S/PDIF, AES-3, IEC-60958, CP-430 receivers, for example the 'last slot' interrupt.

4.3 DIT Transfer Mode

The McASP transmit section may also be configured in digital audio interface transmitter (DIT) mode where it outputs data formatted for transmission over an S/PDIF, AES-3, IEC-60958, or CP-430 standard link. These standards encode the serial data such that the equivalent of 'clock' and 'frame sync' are embedded within the data stream. DIT transfer mode is used as an interconnect between audio components and can transfer multichannel digital audio data over a single optical or coaxial cable.

From an internal DSP standpoint, the McASP operation in DIT transfer mode is similar to the two time slot TDM mode, but the data transmitted is output as a bi-phase mark encoded bit stream with preamble, channel status, user data, validity, and parity automatically inserted into the bit stream by the McASP module. The McASP includes separate validity bits for even/odd subframes and two 384-bit register file modules to hold channel status and user data bits.

If additional serial data pins are used, each McASP may be used to transmit multiple encoded bit streams (one per pin). However, the bit streams will all be synchronized to the same clock and the user data, channel status, and validity information carried by each bit stream will be the same for all bit streams transmitted by the same McASP module.

The McASP can also automatically re-align the data as processed by the DSP (any format on a nibble boundary) in DIT mode; reducing the amount of bit manipulation that the DSP must perform and simplifying software architecture.

4.4 McASP clock generators

The McASP transmit and receive clock generators are identical. Each clock generator can accept a high-frequency master clock input. The transmit and receive bit clocks can also be sourced externally or can be sourced internally by dividing down the high-frequency master clock input (programmable factor /1, /2, /3, ... /4096). The polarity of each bit clock is individually programmable.

A typical usage for the frame sync pins is to carry the left-right clock (LRCLK) signal when transmitting and receiving stereo data. The frame sync signals are individually programmable for either internal or external generation, either bit or slot length, and either rising or falling edge polarity.

Some examples of the things that a system designer can use the McASP clocking flexibility for are:

- Input a high-frequency master clock (for example, 512fs of the receiver), receive with an internally generated bit clock ratio of /8, while transmitting with an internally generated bit clock ratio of /4 or /2. (An example application would be to receive data from a DVD at 48 kHz but output up-sampled or decoded audio at 96 kHz or 192 kHz.)

- Transmit/receive data based one sample rate (for example, 44.1 kHz) using McASP0 while transmitting and receiving at a different sample rate (for example, 48 kHz) on McASP1.

- Use the DSP's on-board AUXCLK to supply the system clock when the input source is an A/D converter.

4.5 McASP Error Handling and Management

To support the design of a robust audio system, the McASP module includes error-checking capability for the serial protocol, data underrun, and data overrun. In addition, each McASP includes a timer that continually measures the high-frequency master clock every 32-SYSCLK2 clock cycles. The timer value can be read to get a measurement of the high-frequency master clock frequency and has a min-max range setting that can raise an error flag if the high-frequency master clock goes out of a specified range.

Upon the detection of any one or more of the above errors (software selectable), or the assertion of the AMUTE_IN pin, the AMUTE output pin may be asserted to a high or low level (selectable) to immediately mute the audio output. In addition, an interrupt may be generated if enabled based on any one or more of the error sources.

4.6 McASP Summary

The two McASPs on the TMS3206713 provide a total of 16 serial lines, independently programmable as transmit or receive. Each McASP has highly flexible independent clock and frame control for its receive and transmit group. Each serial line in turn supports multichannels of TDM data or alternatively direct interface to a variety of digital serial audio data transfer standards. The McASP enables a variety of serial audio interfaces needed in the breadth of high-performance multichannel audio applications.

5 Conclusion

The TMS320C6713 peripheral set enables the device to directly interface to a variety of components in these systems. The McASPs provide highly-flexible direct interconnect to the digital audio streams as well as high performance audio data converters. The two-level cache enables efficient data management and real time I/O while hiding performance issues associated with low cost external SDRAM. The TMS320C6713 DSP device architecture is ideally suited for multichannel, high-performance audio applications.

6   References

1. *TMS320C6713 Floating-Point Digital Signal Processor* data sheet (SPRS186)
2. *TMS320C6211Cache Analysis* application report (SPRA472)
3. *TMS320C6000 DSP Multichannel Audio Serial Port (McASP) Reference Guide* (SPRU041)
4. *TMS320C621x/C671x Two-Level Internal Memory Reference Guide* (SPRU609)
5. *TMS320C6000 CPU and Instruction Set Reference Guide* (SPRU189)
6. *TMS320C6000 Peripherals Reference Guide* (SPRU190)
7. Payan, Reimi, *DSP software and hardware trade-offs in Professional Audio Applications*, Audio Engineering Society, 112th Convention. 2002 May 10–13 Munich, Germany

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

Products		**Applications**	
Amplifiers	amplifier.ti.com	Audio	www.ti.com/audio
Data Converters	dataconverter.ti.com	Automotive	www.ti.com/automotive
DSP	dsp.ti.com	Broadband	www.ti.com/broadband
Interface	interface.ti.com	Digital Control	www.ti.com/digitalcontrol
Logic	logic.ti.com	Military	www.ti.com/military
Power Mgmt	power.ti.com	Optical Networking	www.ti.com/opticalnetwork
Microcontrollers	microcontroller.ti.com	Security	www.ti.com/security
		Telephony	www.ti.com/telephony
		Video & Imaging	www.ti.com/video
		Wireless	www.ti.com/wireless

Mailing Address: Texas Instruments

Post Office Box 655303 Dallas, Texas 75265