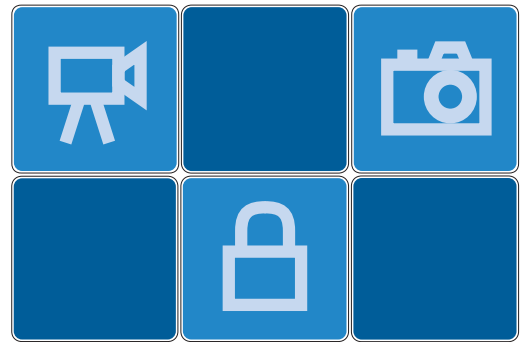




VISIBLE IMPACT.



Silicon IP Cores


2014
Product Catalog



ALMA TECHNOLOGIES®
SILICON IP CORES
2014 PRODUCT CATALOG

Copyright © 2014 ALMA TECHNOLOGIES S.A. All rights reserved.

XILINX and ARTIX are registered trademarks of Xilinx, Inc.
ALTERA, Avalon-ST and CYCLONE are registered trademarks of Altera Corporation.
All other trademarks and trade names are the property of their respective owners.




H.264 Advanced Video Coding (AVC) is an in-force ITU-T recommendation for achieving increased compression efficiency in video coding and has been leading video applications to a new era over the last years.

Alma Technologies has invested a great effort to develop expertise and produce among the finest hardware H.264 implementations available today as silicon IP cores. Fine-tuned for optimum performance, both in terms of pixel throughput and image quality, the H.264 cores perform exceptionally well in a variety of demanding video content situations.

Our H.264 cores are remarkably easy to operate, ready to encode video right out of the box, with simple and comprehensive programming and input/output interfacing. The H.264 cores are designed for easy integration in customer designs and our own or a third party external memory controller may be used equally easy and effectively.

We offer the H.264 encoders in three main configurations: the H264-BP-E core, which is a Baseline profile encoder, the H264-MP-E which is a Main profile encoder and the H264-HP-E which is a High profile encoder. All configurations are also available as All-Intra encoders. All cores achieve high performance and uncompromised video quality. The Baseline core is oriented for use in applications where area or very high output bitrates is a critical factor. The Main Profile core is a bit larger but achieves better compression. The High Profile encoder provides exceptional encoding results, supporting up to 12 bits per sample as well as the 4:2:2 chroma sub-sampling format.

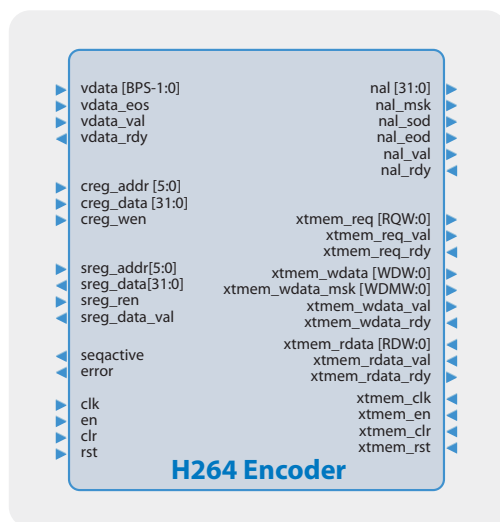


H.264 Encoding IP Cores

The H.264 Encoding family of our silicon Intellectual Property Cores offers standalone and full hardware implementations of the ITU-T H.264 specification. The family includes Baseline, Main and High profile encoders for real time encoding of video streams up to Level 5.2.

Our H.264 encoders require minimal host intervention as they only need to be programmed once per video sequence. Once programmed, an arbitrary number of video frames can be encoded without needing any further intervention or assistance from the host system CPU (CPU-less operation).

Ease of Use

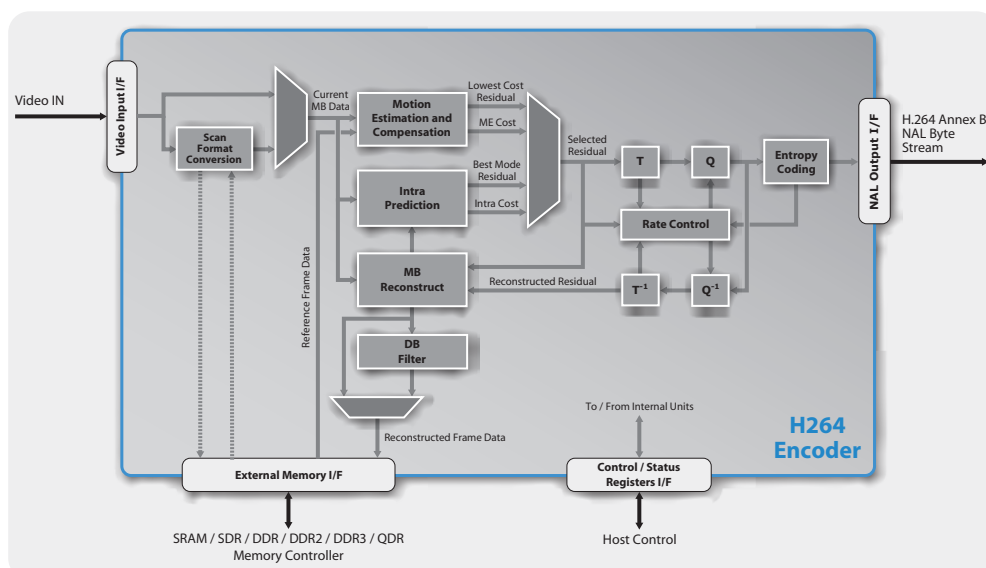


The H.264 encoders use very simple interfaces to communicate with the external host system logic and are setup by means of a simple control register set. The default values for most parameters may be pre-configured to custom specifications. Input and output data interfaces support continuous streaming while also allowing the external logic to fully control the flow of data. A generic memory controller interface is used, supporting a variety of available memory controllers.

Compressing video is as simple as that: connect an external memory controller with the required amount of memory and then program the source video parameters - if different than the default values - using the control registers of the core. Following that, the core is ready to accept and compress the input video stream, producing a compliant H.264 Annex B NAL Byte stream on the output. The external application may monitor the core's status by probing the status registers of the encoder.

High Performance

The H.264 Encoding IP Cores are designed to achieve high throughput with lower clock frequencies. At the same time, a big effort was made to keep the core's silicon area requirements limited so that the core can fit and process Full HD video even in low-end FPGA devices. As an example, the core is capable of encoding Full HD 1080p30 video on the Altera Cyclone III or on the Xilinx Artix-7 devices. In ASIC implementations the core is capable to exceed the 120 fps mark for Full HD video.



Specifications

The H.264 Encoding IP Cores can accept the uncompressed input video stream in planar, interleaved, or macroblock scan format. The output is standard compliant Annex B NAL Byte stream format. No post processing on the output stream, other than saving or transmitting, is required by the host. The output NAL Byte stream can be decoded, as is, by any ITU-T H.264 compliant decoder that satisfies the level requirements of the stream and conforms to the respective ITU-T H.264 profile.

The H.264 encoding cores implement a simple – yet flexible – request based external memory interface with independent read and write data paths. This makes the cores independent of memory type, supporting for example operation with SRAM, SDRAM, DDR, DDR2 or DDR3 types of memory. Glue-less connected external memory controllers are also available. The encoders are designed to be tolerant to memory delays or latencies, which may be present on shared memory system architectures.

Features

Advanced H.264 Implementation

- High throughput implementation: Sustained 2.5 clock cycles per pixel worst case processing rate.
- Superior compression and video quality from QCIF to HD resolutions.
- CQP - VBR encoding mode:
 - Rate-Distortion optimized output.
 - Up to 240 MBits/s output.
- CBR encoding mode:
 - HRD CPB compliant CBR NAL output.
 - Sub-frame operation with tunable number of macroblocks basis.
 - Further micro adjustment of quantization per macroblock maximizes the perceived video quality.
 - Based on run-time continuous self-trained models for excellent adaptability to spatial and temporal video variations.
 - Tunable and independent to video GOP size operation, enables video system latency control with respect to initial buffering requirements.
 - On-the-fly rate changes are supported.
 - Up to 240 MBits/s output.
- Full search, variable block size, sub-pixel motion estimation engine:
 - Significantly more immune to video content than partial searches.
 - Eliminates local minima traps.
 - 32x20 or higher search area around the partition origin.
 - Up to 4 motion vectors per macroblock.
 - Full, half and quarter pixel accuracy.
 - Single reference frame.
 - Sophisticated block skipping logic enables advanced low bit rate encoding with minimized motion artifacts.
- Advanced Intra prediction: All prediction modes supported.
- Error resilience.
- Optional advanced thresholding of quantized transform coefficients.
- Run-time tunable operation enables decoder compatibility trade-offs.

Smooth System Integration

- Full abstraction of the internal implementation details and the H.264 complexity from the top level I/O and its functionality.
- Simple, microcontroller like, programming interface.
- High speed, flow controllable, streaming I/O data interfaces:
 - Simple and FIFO like.
 - Full data flow control without extra clock cycles penalty.
 - Avalon-ST™ compliant (ready latency 0).
- Registered I/O ports.
- Low native encoding latency:
 - In macroblock scan video input the encoding latency is approximately 1280 video pixels.
 - In interleaved scan video input the encoding latency is approximately 16 video lines.
 - In planar scan video input the encoding latency is approximately one frame.
- Low requirements in external memory bandwidth.
- Flexible external memory interface:
 - Independent of external memory type.
 - Tolerant to latencies.
 - Allows for shared memory access.
 - Can optionally operate on independent clock domain.



H.264 Encoding IP Cores Selection Matrix

|  | H264-BP-E | H264-BPI-E | H264-MP-E | H264-MPI-E | H264-HP-E | H264-HPI-E |
|---|----------------------|----------------------|-----------|------------|--|---|
| Profile | Constrained Baseline | Constrained Baseline | Main | Main | High 10 High 4:2:2 High 4:4:4 Predictive | High 10 Intra High 4:2:2 Intra High 4:4:4 Intra |
| Chroma format(s) | 4:2:0 | 4:2:0 | 4:2:0 | 4:2:0 | 4:2:0, 4:2:2 | 4:2:0, 4:2:2 |
| Sample depth (bits) | 8 | 8 | 8 | 8 | 8, 10, 12 | 8, 10, 12 |
| Slice type(s) | IDR, P | IDR | IDR, P | IDR | IDR, P | IDR |
| Multiple slices encoding option | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| CAVLC | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| CABAC | – | – | ✓ | ✓ | ✓ | ✓ |
| Rate control (CBR & CQP-VBR) | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Separate Luma / Chroma QP control | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Ultra Low sub-frame transmission latency encoding | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Annex B NAL Byte stream compliant output | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| AVC-Intra 50 / 100 / Ultra suitability | – | – | – | – | ✓ | ✓ |
| Standalone CPU-less operation | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Throughput (clocks/pixel) | 2.5 | 2.5 | 2.5 | 2.5 | 2.5 - 2.75 | 2.5 - 2.75 |
| Silicon requirements | high | medium | high | medium | high | medium |
| Available for ASICs | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Altera FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Lattice FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Xilinx FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |


 : Supported
 – : Not Applicable



A complete series of IP Cores covering almost every application requiring image compression is available.

All our still image compression IP Cores have been carefully crafted to offer maximum image quality and performance, while maintaining efficiency and simplicity during the integration.

Still image compression algorithms are equally important for video compression applications. Applications sensitive to latency, throughput, power or silicon area, can benefit from the efficiency of the still image compression algorithms.



Still Image Compression IP Cores

Alma Technologies offers a complete family of still image compression IP Cores, covering almost every application requiring image compression. Most of the well-established still image compression algorithms are implemented, such as the JPEG, the JPEG 2000, Lossless JPEG and JPEG-LS.

The following selection matrix can help in deciding which still image compression scheme is best suited for each application.


Still Image Compression Selection Matrix

| | Baseline JPEG | Extended JPEG | Lossless JPEG | JPEG-LS | JPEG 2000 |
|---------------------------------------|---------------|---------------|----------------|----------------|-----------|
| Lossy compression | ✓ | ✓ | ✗ ¹ | ✗ ¹ | ✓ |
| Numerically lossless compression | ✗ | ✗ | ✓ | ✓ | ✓ |
| Lossy compression efficiency | very good | very good | – | – | excellent |
| Lossless compression efficiency | – | – | good | excellent | excellent |
| Maximum bits per sample | 8 | 12 | 16 | 16 | 16 |
| Grayscale | ✓ | ✓ | ✓ | ✓ | ✓ |
| Color | ✓ | ✓ | ✓ | ✓ | ✓ |
| Rate control | ● | ● | ✗ | ✗ | ✓ |
| Multiple quality layers | ✗ | ✗ | ✗ | ✗ | ● |
| Region of interest | ✗ | ✗ | ✗ | ✗ | ● |
| Standalone CPU-less operation | ✓ | ✓ | ✓ | ✓ | ✓ |
| Requires external memory ² | no | no | no | no | yes |
| Silicon requirements | low | low | very low | low | high |
| Available for ASICs | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Altera FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Lattice FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Microsemi FPGAs | ✓ | ✓ | ✓ | ✓ | ✗ |
| Available for Xilinx FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ |

Notes :

- ✓ : Supported
- : Optional
- ✗ : Not Supported

¹ Near-lossless function(s) available.
² Typical use case.



Baseline and Extended JPEG IP Cores

The JPEG cores provide an efficient and low cost image compression / decompression solution based on the widely adopted JPEG image compression standard.

A full range of JPEG IP core solutions is available, combining high performance with area efficient architectures and support for Motion JPEG streaming with Rate Control. The JPEG cores are standalone; no extra software pre-processing or post-processing is required. Additional functions that may be needed, such as Color Space Conversion or Raster to Block Scan Converter, are also available.

Features

Baseline ISO/IEC 10918-1 JPEG Compliance

- Programmable Huffman Tables (two DC, two AC).
- Programmable Quantization Tables (up to four).
- Up to four color components.
- Supports all possible scan configurations and all JPEG formats for input and output data.
- Supports any image size up to 64K x 64K.
- Supports DNL and restart markers.

Additional Processing Capabilities

- Motion JPEG payload support.
- Video mode Rate-Control engine (optional).

Ease of Integration

- Registered I/O ports.
- Simple, microcontroller like, programming interface.
- High speed, flow controllable, streaming I/O data interfaces:
 - Simple and FIFO like.
 - Avalon-ST™ compliant (ready latency 0).
- Decoding:
 - Stand alone operation.
 - Automatic self-programming by JPEG markers parsing.
 - Marker errors catching.
 - Broadcasting of decoded image parameters for controlling peripherals.
- Encoding:
 - Single clock per input sample processing rate.
 - Fully programmable through standard JPEG marker segments.
 - Automatic JPEG markers generation on the output.
 - Automatic program-once encode-many operation.

Our JPEG implementations are fully compliant with the Baseline or the Extended Sequential DCT mode of the ISO/IEC 10918-1 JPEG standard. This makes the JPEG cores ideal for interoperable systems and devices such as digital cameras, camcorders, office automation equipment, medical imaging systems, video conference systems and remote surveillance systems.

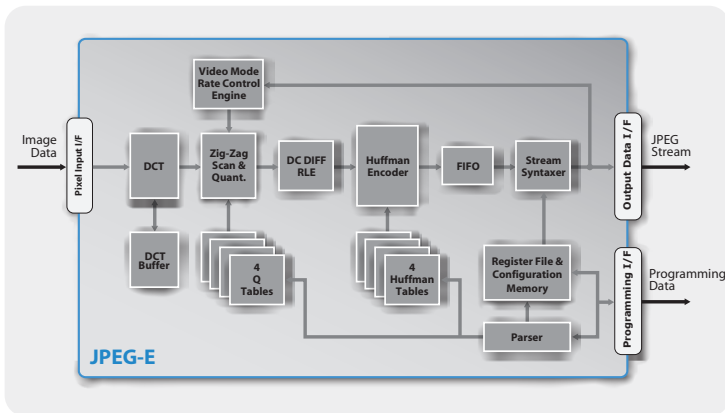
In addition to generating standalone JPEG streams, the JPEG encoding cores are also capable of producing the (de facto) standard video payload of many Motion JPEG container formats (e.g. USB Video Class devices). Furthermore, bandwidth constrained applications will benefit from our unique and advanced Rate Control engine.

All cores feature easy-to-use, fully controllable and FIFO-like streaming input and output interfaces. Being carefully designed, rigorously verified and multiple times silicon-proven, the JPEG cores offer a reliable and easy to integrate solution.



JPEG-E, JPEG-E-X

Baseline JPEG Encoder, Extended JPEG Encoder



The JPEG-E and JPEG-E-X cores are standalone and high-performance JPEG encoders for still image and video compression applications.

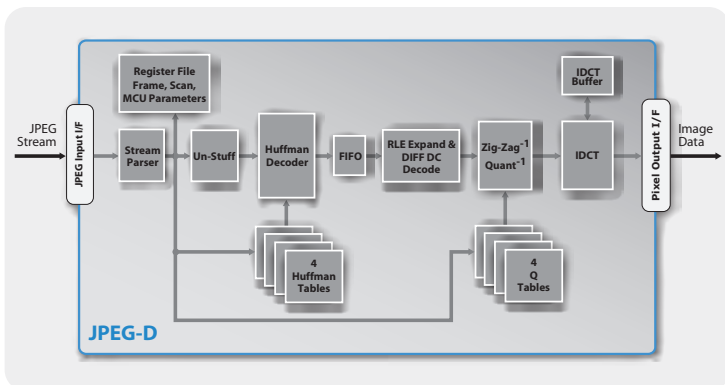
Among the fastest available JPEG encoders, the cores are able to encode at Full HD (1080p30) or higher rates, even in low-end FPGA devices. The JPEG-E-X encoder additionally supports up to 12 bits sample depth, making it well suited for high dynamic range applications such as professional video, machine vision and medical applications.

The cores include optionally an advanced Rate Control engine which provides important flexibility for video applications with constrained bandwidth specifications.



JPEG-D, JPEG-D-X

Baseline JPEG Decoder, Extended JPEG Decoder



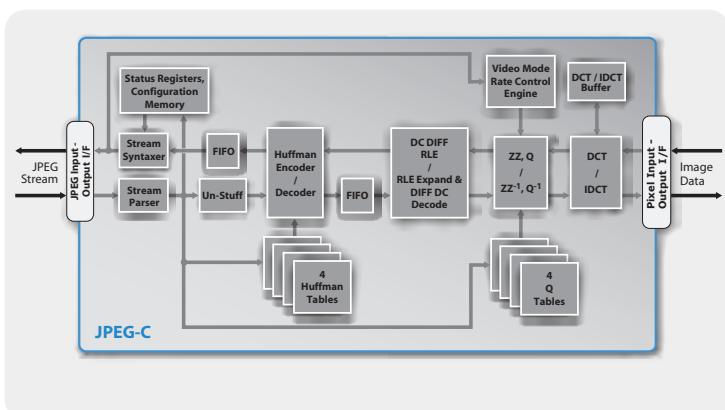
The JPEG-D and JPEG-D-X cores are standalone and high performance JPEG decoders for still image and video decompression applications.

Among the fastest available JPEG decoders, the JPEG-D / JPEG-D-X cores can decode at Full HD (1080p30) or higher rates, even in FPGA devices. The JPEG-D-X decoder additionally supports up to 12 bits sample depth decoding.



JPEG-C

Baseline JPEG Codec



The JPEG-C core is a standalone and high performance, half-duplex Baseline JPEG codec for still image and video compression applications.

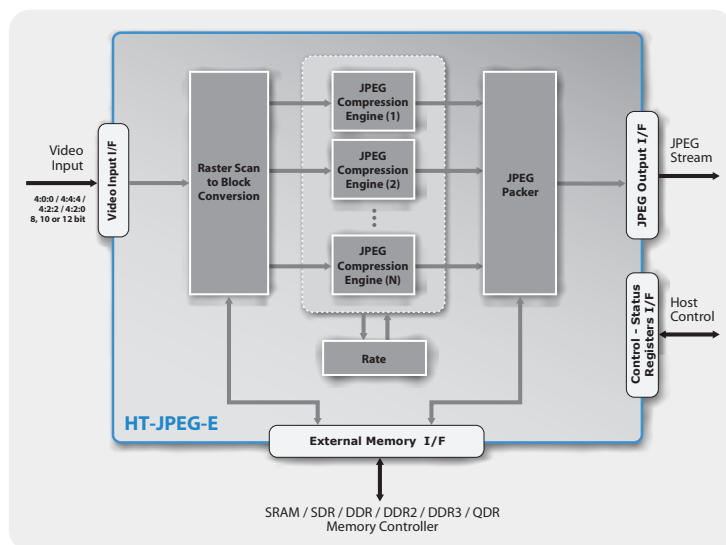
One of the fastest available JPEG codecs, the JPEG-C can operate at Full HD (1080p30) or higher rates, even in FPGA devices.

The core includes optionally an advanced Rate Control engine which provides important flexibility for video applications with constrained bandwidth specifications.



HT-JPEG-E

High Throughput Baseline and Extended JPEG Encoder




The HT-JPEG-E core is a standalone and high-performance JPEG encoder designed for applications that require ultra high frame rate SD and HD encoding and Ultra HD (4K/8K and beyond) video encoding, even in low-end ASIC or FPGA silicon. The core implements an advanced Baseline / Extended Sequential DCT ISO/IEC 10918-1 JPEG compliant encoder. Real time encoding of 4:0:0, 4:4:4, 4:2:2 and 4:2:0 video streams are supported, in 8, 10 or 12 bits sample depths.

The HT-JPEG-E core requires minimal host intervention as it only needs to be programmed once per video sequence. Once programmed, it can encode an arbitrary number of video frames without the need of any further actions by the host system CPU.

The core includes an advanced Constant Bitrate Control engine which provides important flexibility for video applications with constrained bandwidth specifications.

The HT-JPEG-E core accepts the uncompressed raw video data in interleaved scan format. It outputs standalone, standard compliant, JPEG byte stream format. No post processing on the output stream is required by the host. The output JPEG byte stream can be decoded, as is, by any corresponding ISO/IEC 10918-1 JPEG compliant decoder.



Lossless JPEG IP Cores

We offer Lossless JPEG coding solutions suitable for applications that need to compress images and reproduce them without any loss of quality. In fact, reproduced images are bit by bit identical to the original uncompressed images. The Lossless JPEG IP Cores are ideal for medical imaging applications, high-end digital photography or professional video equipment applications.

The Lossless coding mode of JPEG is an efficient light image compression algorithm which achieves good compression results, without the complexity overhead of more sophisticated solutions such as the JPEG 2000. Lossless JPEG was added to the ITU-T JPEG recommendations in 1995 (ISO/IEC 10918-1 standard, ITU T.81 recommendation).

Features

ISO/IEC 10918-1 Compliance

- Conforms to the spatial (sequential) lossless encoding mode (SOF3), of the ISO/IEC 10918-1 standard (CCITT T.81 recommendation).
- Standalone operation:
 - Pixel samples input / output.
 - Standalone - ISO/IEC 10918-1 compliant - JPEG stream input / output.
- Easily programmable through standard JPEG markers stream.
 - Programmable image dimensions.
 - Programmable image sample precision (2 - 16 bits/sample).
 - Up to 4 programmable Huffman tables.
 - Programmable Restart Interval.
 - Programmable Point Transform function.
 - Programmable APPn and COM markers.
- Programming errors catch-up features.

Limitations

- Up to three image components are supported (Nf field of the SOF3 marker segment = 1 or 2 or 3).
- Single scan encoding (Only one SOS marker segment, with Ns field = Nf).
- No DNL marker insertion (Y field of the SOF3 marker segment > 0).
- Fixed parameters:
 - No sub-sampling (Hi and Vi fields of the SOF3 marker segment = 1).
 - Prediction function is fixed to left-hand predictor (predictor 1) (Ss field of SOS marker segment = 1).

The LJPEG cores implement Lossless JPEG compression / decompression in a compact, high-performance, stand-alone package ideal for applications where bit-by-bit accurate reproduction of an image is essential.

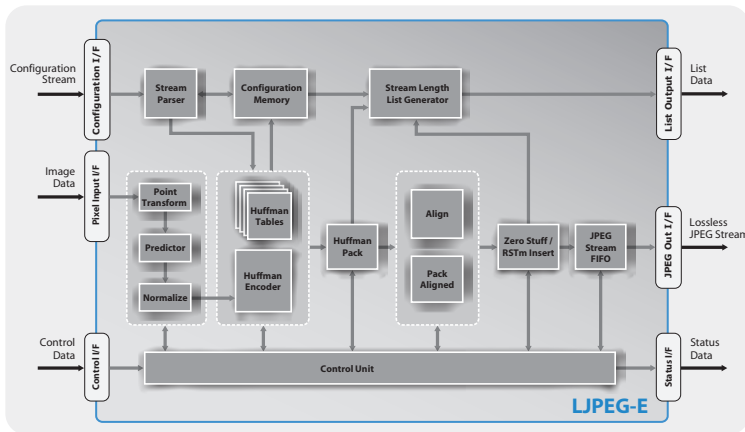
The LJPEG cores conform to the spatial (sequential) lossless encoding mode (SOF3) of the ISO/IEC 10918-1 standard (ITU T.81 recommendation). Rather than the Discrete Cosine Transform (DCT) functions used in the lossy JPEG compression - which can introduce round-off errors - the LJPEG cores employ a reversible predictor function as described in the specification. Encoding and decoding images with no information loss is thus possible, while also requiring a significantly smaller silicon footprint than the standard - lossy JPEG implementations.

The cores are full, standalone, CPU-less solutions and feature easy-to-use fully controllable and FIFO-like streaming input and output interfaces. Being carefully designed, rigorously verified and silicon-proven, the LJPEG cores offer a reliable and easy to integrate solution.



LJPEG-E

Lossless JPEG Encoder



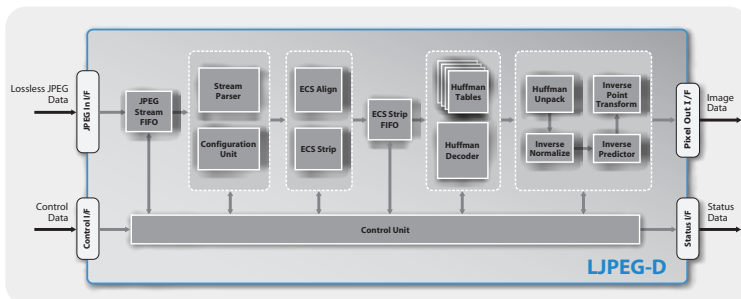
The LJPEG-E core implements the Lossless JPEG compression in a compact, high-performance, stand-alone package, ideal for applications where bit-by-bit accurate reproduction of an image is essential.

Designs show that the core requires just 21K equivalent NAND2 gates in an ASIC (90 nm process under typical conditions) and that it also fits in a variety of low-end and high-end FPGA devices. Its heavily optimized architecture enables also a very high throughput, reaching 500 MSamples/s on ASICs and 200 MSamples/s on FPGA devices.




LJPEG-D

Lossless JPEG Decoder



The LJPEG-D core implements a Lossless JPEG decoder in a compact, high-performance, stand-alone package, ideal for applications where bit-by-bit accurate reproduction of an image is essential.

Designs show that the core requires just 36K equivalent NAND2 gates in an ASIC (90 nm process under typical conditions) and that it also fits in a variety of low-end and high-end FPGA devices. Its optimized architecture enables also a high throughput, reaching 250 MSamples/s on ASICs and 95 MSamples/s on FPGA devices.



JPEG-LS IP Cores

We offer JPEG-LS coding solutions suitable for applications that need numerically lossless or near-lossless image compression, such as professional video equipment, machine vision, medical and satellite imaging devices. The JPEG-LS is an excellent lossless image compression algorithm. It achieves results highly competitive to those offered by far more complex solutions, such as the JPEG 2000. The JPEG-LS cores are compliant to the ISO/IEC 14495-1 standard.

Features

ISO/IEC 14495-1 JPEG-LS Compliance

- Programmable image dimensions from 8 x 8 up to 64K x 64K.
- Grayscale, 4:4:4, 4:2:2, 4:1:1 and 4:2:0 chroma subsampling formats.
- Programmable sample precision (BPP) from 2-bits up to 16-bits.
- Programmable point transform (up to BPP-1).
- Programmable local gradient thresholds (up to 2x BPP-1) and context parameters reset threshold value (up to 127).
- Header error catch-up features.

Ease of Integration

- Registered I/O ports.
- Simple, microcontroller like, programming interface.
- High speed, flow controllable, streaming I/O data interfaces:
 - Simple and FIFO like.
 - Avalon-ST™ compliant (ready latency 0).
- Single clock per input sample processing rate.
- Fully programmable through standard JPEG-LS marker segments.
- Automatic JPEG-LS markers generation on the output.
- Automatic program-once encode-many operation.

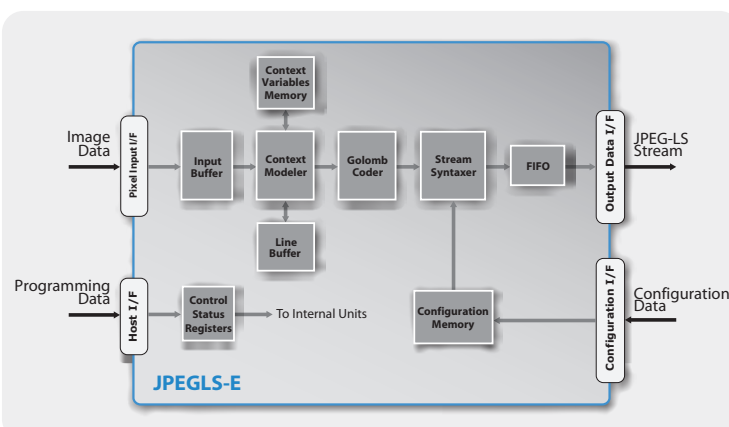
JPEG-LS was developed to provide a low complexity lossless image compression specification with better compression potential than Lossless JPEG. The algorithm at the core of JPEG-LS is LOCO-I (Low Complexity - Lossless Compression for Images). It uses a non-linear predictive scheme with rudimentary edge detecting capability, based on the four nearest - causal - neighbors (left, upper left, upper and upper right) and an entropy encoder which uses adaptively selective Golomb-type codes. The low complexity scheme of JPEG-LS is based on the assumption that prediction residuals follow a two-sided geometric distribution and the fact that Golomb-codes are optimal for geometric distributions, thus the modeling and coding units are matching.

The cores are full, standalone, CPU-less solutions and feature easy-to-use fully controllable and FIFO-like streaming input and output interfaces. Being carefully designed, rigorously verified and silicon-proven, the JPEG-LS cores offer a reliable and easy to integrate solution.




JPEGLS-E

JPEG-LS Encoder



The JPEGLS-E core is a JPEG-LS encoder that forms a high performance solution for lossless image and video compression applications. The JPEGLS-E can encode at Full HD (1080p30) or higher rates, even in FPGA devices. Full compliance to the ISO/IEC 14495-1 standard makes the JPEG-LS encoder core ideal for any open platform application where interoperability is a critical factor.

The JPEGLS-E core is available in two configurations: The true numerically lossless only encoding configuration offers the highest throughput, while the numerically-lossless and near-lossless encoding configuration offers the maximum application flexibility.



JPEG 2000 IP Cores

Alma Technologies was one of the first companies worldwide to offer a hardware JPEG 2000 IP Core solution back in 2002. Today, the latest generation of our JPEG 2000 encoder is one of the more capable and flexible JPEG 2000 solutions available. Several objectives were set and achieved during the core's design: high performance, scalability, robustness and ease of integration.

Features

ISO/IEC 15444-1 JPEG 2000 Image Coding System compliance

- Grayscale and three or four component color images.
- 4:4:4, 4:2:2, 4:1:1 and 4:2:0 chroma subsampling formats.
- Up to 16-bits per component sample precision.
- Up to 65535 x 65535 image resolution.
- Up to 8192 x 8192 tile resolution.
- Lossless or lossy compression.
- Advanced rate control engine.
- Single or multiple quality layers encoding.
- CPRL progression order.
- LRCP progression order (grayscale only).
- Error resilient encoding features.
- Standard compliant code stream (JPC) or file (JP2) output.

Programmable JPEG 2000 Encoding Options

- Image and pixel input format (frame and tile size, number of components, pixel depth, subsampling format and scan order).
- Wavelet filter type (5/3 or 9/7) and number of wavelet transform levels.
- Code-block size (64 or 32 or 16 on each dimension).
- Quantization tables.
- Entropy coding switches (reset, restart, segmark).
- Number of quality layers (grayscale only, up to 30).
- Output bitrate (per quality layer).
- Progression order and JPEG 2000 output format (proprietary or JPC or JP2).
- Operation on an entire image, or on tile-component, or on tile basis.

Ease of Integration

- Registered I/O ports.
- Simple, microcontroller like, programming interface.
- High speed, flow controllable, streaming I/O data interfaces:
 - Simple and FIFO like.
 - Full data flow control without extra clock cycles penalty.
 - Avalon-ST™ compliant (ready latency 0).
- Single port external memory interface compatible with most (multi-port) memory controllers.
- Tunable design, enables silicon, frequency and throughput trade offs.
- Flexible external memory interface:
 - Independent of external memory type.
 - Glue-less connection to external memory controllers.
 - Tolerant to latencies.
 - Allows for shared memory access.
 - Can optionally operate on independent clock domain.

Full compliance to the ISO/IEC 15444-1 JPEG 2000 standard makes the JPEG2K-E core ideal for interoperable systems and devices. The JPEG2K-E supports up to 8Kx8K image or tile resolution and it can sustain a very high, over 1080p30, throughput. The JPEG2K-E includes also an advanced post-compression, rate-distortion optimized, rate control engine which provides full control over the required bandwidth on the JPEG 2000 stream output. The bitrate of the JPEG 2000 stream can be accurately adjusted while, at the same time, preserving the maximum image fidelity that is possible within the available bandwidth constraints.

The JPEG2K-E core implements a simple but yet flexible, requests based, external memory interface with independent read and write data paths. This makes the JPEG2K-E independent of memory type supporting, for example, operation with SRAM, SDRAM, DDR, DDR2 and DDR3 types of memory. Glue-less connected external memory controllers are also available. JPEG2K-E is designed to be tolerant to memory delays and latencies, which may be present on shared memory system architectures.

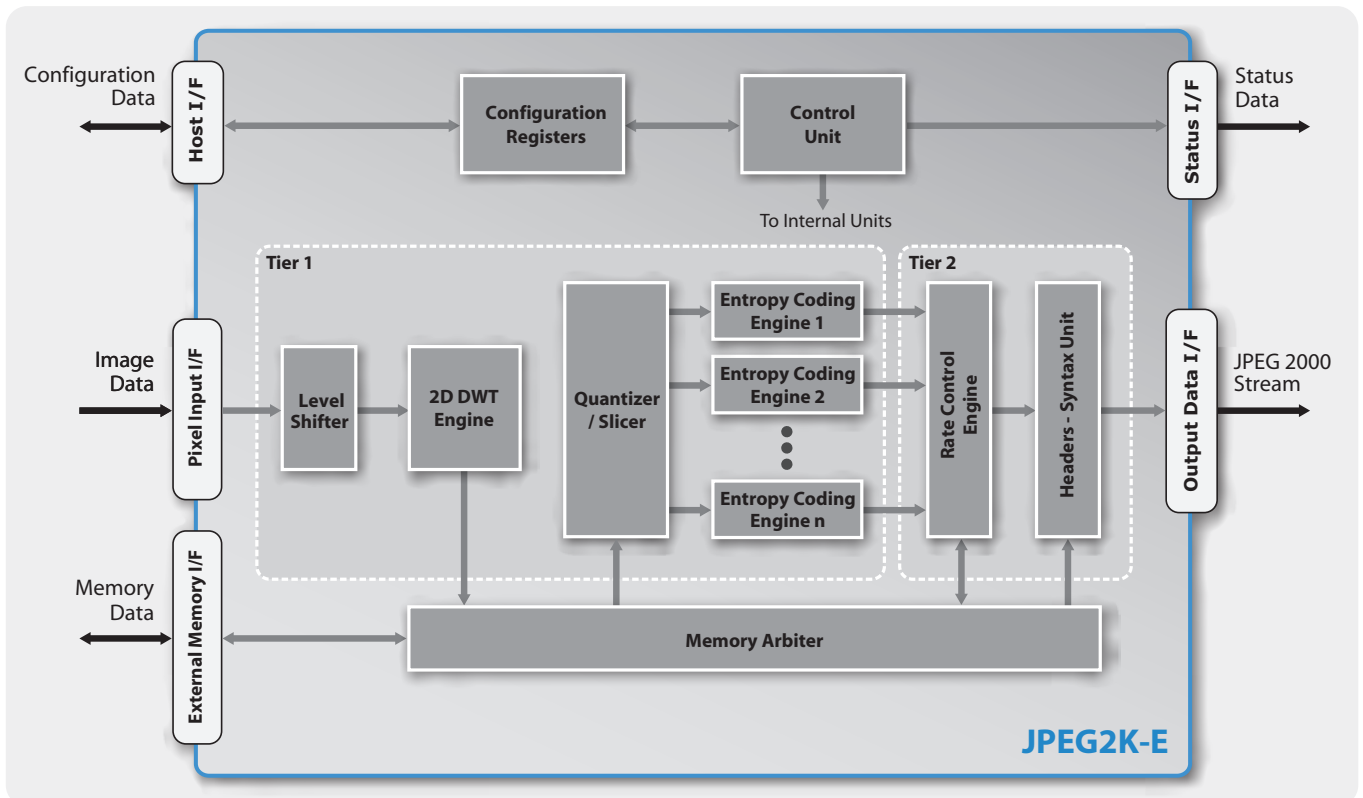
The core is designed with easy to use, fully controllable and FIFO-like, streaming input and output interfaces. Being carefully designed, rigorously verified and silicon-proven, the JPEG2K-E is a reliable and easy to integrate core.




JPEG2K-E

JPEG 2000 Encoder

The JPEG2K-E core is a video and still image encoder that implements the JPEG 2000 lossy and lossless image compression standard. The JPEG 2000 standard offers an advanced quality and feature set, lending itself to a wide range of uses from digital cameras through to space imaging and other key sectors.





Alma Technologies offers a series of cryptographic functions, suitable for encrypting or decrypting a data stream or message, as well as authenticating the encrypted or decrypted data. Highly efficient implementations of Advanced Encryption Standard (AES) are included, and implementations of Hash functions such as the SHA256, SHA-1 and MD5.

The implementation of the AES is highly flexible and also implements the most commonly used Block Cipher Modes of operation, ECB, CBC, CFB, OFB, CTR and the GCM Authenticated Encryption / Decryption mode.



AES Block Cipher IP Cores

The product family of our AES IP cores includes implementations of the most commonly used Block Cipher modes of operation in two different bitwidths which affect the footprint and the throughput of the solution.

The following selection matrix illustrates what is featured in each IP core variation and can be used as a guide to select the best suited product for a given application.

AES Products Selection Matrix

| | AES32-C | AES128-C | AES32-P | AES128-P | AES-GCM128 |
|--|-----------------------|-----------------------|-----------------------|-----------------------|------------|
| Run-time Selection of Cipher-Key length | ✓ | ✓ | ✓ | ✓ | ✓ |
| Run-time Selection of Block Cipher mode | ✗ | ✗ | ✓ | ✓ | – |
| Run-time Selection between Encryption / Decryption | ✓ | ✓ | ✓ | ✓ | ✓ |
| AES Round Key Table included | ✓ | ✓ | ✓ | ✓ | ✓ |
| AES Key Expander included | ● | ● | ● | ● | ✓ |
| ECB mode support | ✓ | ✓ | ✓ | ✓ | – |
| CBC mode support | ✓ | ✓ | ✓ | ✓ | – |
| CFB mode support | ✓ | ✓ | ✓ | ✓ | – |
| OFB mode support | ✓ | ✓ | ✓ | ✓ | – |
| CTR mode support | ✓ | ✓ | ✓ | ✓ | – |
| GCM mode support | ✗ | ✗ | ✗ | ✗ | ✓ |
| Data path bitwidth | 32 | 128 | 32 | 128 | 128 |
| Processing Cycles ¹ | 44/52/60 ² | 11/13/15 ² | 44/52/60 ³ | 11/13/15 ⁴ | 12/14/16 |
| Standalone CPU-less operation | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for ASICs | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Altera FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Lattice FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Microsemi FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ |
| Available for Xilinx FPGAs | ✓ | ✓ | ✓ | ✓ | ✓ |

Notes :

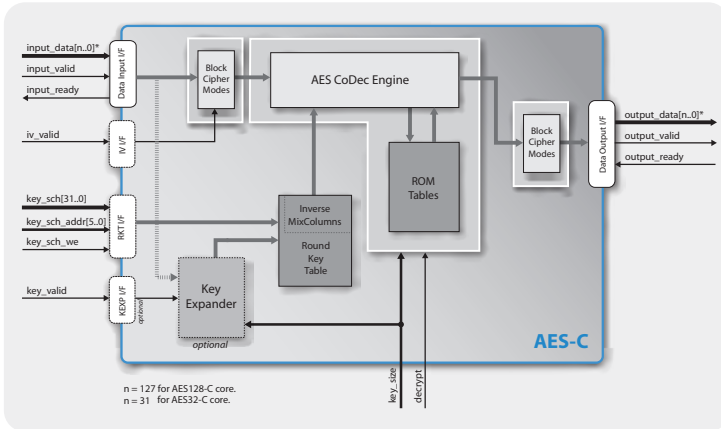
- ¹ Throughput for continuous data input.
- ² An additional cycle is required on CFB, OFB, CTR modes.
- ³ An additional cycle is required on CFB mode.
- ⁴ An additional cycle is required on CBC, CFB, OFB, CTR modes.

- ✓ : Supported
- : Optional
- ✗ : Not Supported
- : Not Applicable



AES32-C, AES128-C

AES Encryption / Decryption, Configurable Block Cipher Mode



The AES-C core family implements the FIPS-197 Advanced Encryption Standard, and can be programmed to either encrypt or decrypt 128-bit blocks of data, with a 128-bit, 192-bit or 256-bit cipher-key.

The AES-C core is available in two variations, standard AES32-C and fast AES128-C. AES32-C has a 32-bit internal data path and the AES128-C has an 128-bit internal data-path. The AES32-C core is more compact in size but with lower throughput than the AES128-C core. For the AES32-C core, 44/52/60 clock cycles are required to encrypt or decrypt an input block with a 128/192/256-bit cipher key respectively. 11/13/15 clock cycles are required for the AES128-C core.

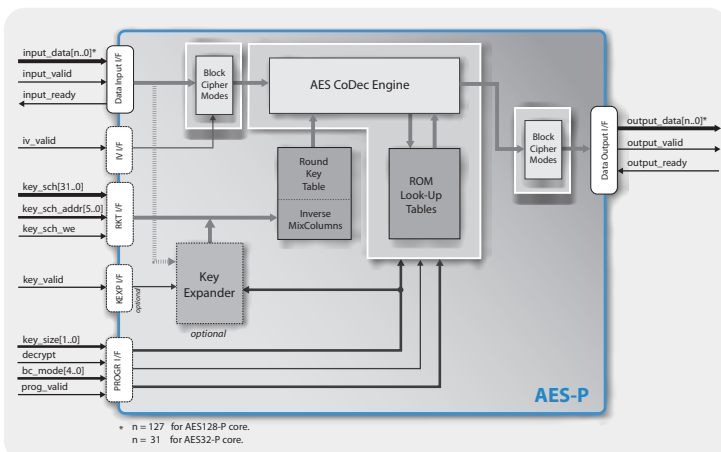
During each step of the encryption / decryption processing, the core requires a previously calculated Round Key Value derived from the cipher-key using a key expansion algorithm. The Round Key Values must be stored to the internal Round Key Table, from which the core acquires the appropriate for each processing step value. Alternatively, instead of directly programming the Round Key Values to the Round Key Table, an optional Key Expander module is provided. The cipher-key is given to the core and the Key Expander module automatically calculates the Round Key Values and fills the internal Round Key Table.

An included configurable wrapper surrounds the AES-C cores, which implements a Block Cipher mode of operation. Block Cipher modes supported are: ECB, CBC, CFB, OFB and CTR.



AES32-P, AES128-P

AES Encryption / Decryption, Programmable Block Cipher Mode



The AES-P core family implements the FIPS-197 Advanced Encryption Standard, and can be programmed to either encrypt or decrypt 128-bit blocks of data, with 128-bit, 192-bit or 256-bit cipher-key. The Block-Cipher mode of operation can also be run-time programmed to one of: ECB, CBC, CFB, OFB, CTR.

The AES-P core is available in two variations, Standard AES32-P and Fast AES128-P. AES32-P has a 32-bit internal data path and the AES128-P has an 128-bit internal data-path.

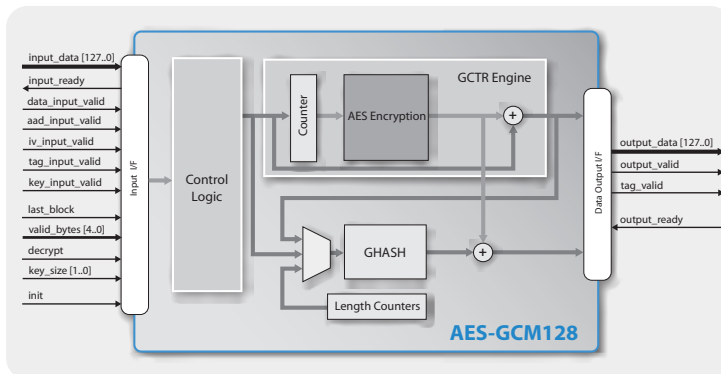
During each step of the encryption / decryption processing, the core requires a previously calculated Round Key Value derived from the cipher-key using a key expansion algorithm. The Round Key Values must be stored to the internal Round Key Table, from which the core acquires the appropriate for each processing step value. Alternatively, instead of directly programming the Round Key Values to the Round Key Table, an optional Key Expander module is provided. The cipher-key is given to the core and the Key Expander module automatically calculates the Round Key Values and fills the internal Round Key Table.

The AES-P cores include run-time programmable Block-Cipher mode selection. Supported modes are: ECB, CBC, CFB, OFB and CTR.




AES-GCM128

GCM-AES, Authenticated Encryption / Decryption



The AES-GCM128 core implements the GCM-AES authenticated encryption / decryption function, as specified in NIST's SP800-38D recommendation for GCM and GMAC, and FIPS-197 Advanced Encryption Standard. The core can be programmed to either encrypt or decrypt 128-bit blocks of data, with a 128-bit, 192-bit or 256-bit cipher key. In addition a Hash value – the Tag – is calculated using the GHASH algorithm for the encrypted or additional plaintext data. In decryption mode, the calculated TAG is compared with the TAG that accompanies the ciphertext data, and a Fail or Pass flag is generated.


The AES-GCM128 core has a 128-bit data path size, meaning 1 clock cycle is required to load/unload each 128-bit plaintext/ciphertext block. A Key Expander is included for the AES in the AES-GCM128 core to automatically generate the Round Key Values for AES. Since the core has a 128-bit datapath and all internal operations are performed on 128-bit words, 12/14/16 clock cycles per HASH are required per 128-bit block. The AES-GCM128 core supports 96-bit Initialization Vectors and input / output Tags of configurable length.



Expanding to the area of Memory Controllers, Alma Technologies currently offers a rich in features SPI Flash Memory Controller, that offers reliable communication up to the very high bitrates supported by the most recent devices.

The SPI Flash Memory Controller is a mature design proven in multiple ASIC designs, and is highly configurable aiming to provide support even for future devices.

Alma Technologies also offers SDRAM memory controllers, supporting SDR, DDR, DDR II memory devices.



SPI Memory Controller

The SPI Memory Controller provides the necessary functionality to a host application in order to communicate with a serial SPI Flash memory device. The controller supports three types of memory accesses: read, write and erase, as well as custom instructions programmable during run-time.

Features

Device Independent

- Automatic identification of a variety of memories.
- Configurable memory features to allow support for more serial flash devices.
- Programmable custom SPI instructions provide support for new SPI Flash device features.

High Performance

- Supports single SPI, dual output SPI, dual input / output SPI, quad input / output SPI.
- Separate Clock domains for SPI subsystem and host data buses.
- Programmable number of dummy cycles on Quad mode Read operations.
- Programmable capture delay for MISO inputs permits large SPI bus round-trip delays.
- Configurable internal FIFOs and programmable thresholds enable full bandwidth utilization.
- Automatic identification of maximum bandwidth access mode among single SPI, dual output SPI, dual input / output SPI, quad input / output SPI.

Flexible Access Model

- Registered Mapped I/O.
- Read access sizes from 1 byte up to memory density.
- Read accesses starting from any address offset.
- Write access sizes from 4 bytes up to memory density.
- Write accesses starting from any address offset that is multiple of 4.
- Erasure of:
 - any sector (4KB),
 - any block (64KB),
 - whole chip.

Ease of Integration

- Auto-detection of a wide set of serial flash devices to minimize programming overhead.
- Auto detection of the fastest way to read or program the memory, to maximize bandwidth and minimize programming overhead.
- Deep Power-down Mode support to minimize power consumption.
- Optional Execute-on-the-Fly™ interface for better CPU connectivity.
- Optional Block Read interface that can transfer a block of data with minimal programming.
 - Block Read can be automatically initiated after Power-On, or on demand during run-time.

A host can interface to the Serial Flash in a number of ways. Transferring data from the Flash memory to the host is done with minimum effort with a Block Read Interface, which uses a DMA mechanism to transfer a block of data to the host's memory space. Alternatively, the host can introduce a read request using the core's programmable registers. Then the core serves this request and sends the necessary instructions to the Serial Flash device. An additional RAM like interface which permits on-the-fly code execution is also available.

The design uses two clock domains. The first clock domain is used for the host interface of the core, while the second clock domain is used for generating the SPI clock and must be at least 2 times faster than the target SPI clock.

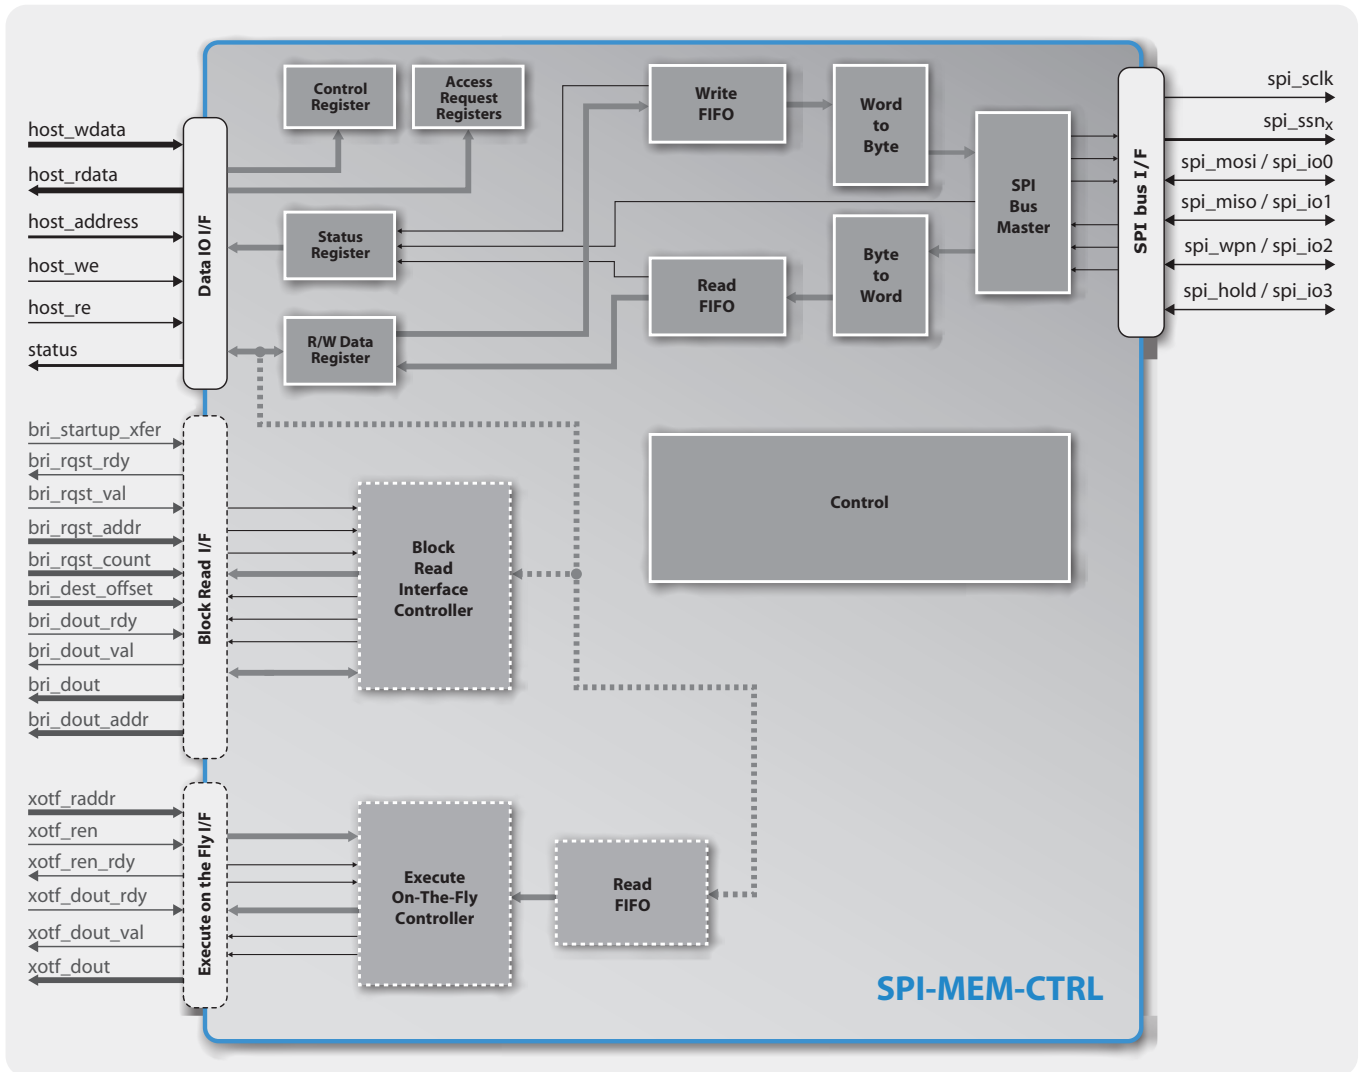
The SPI-MEM-CTRL core features configurable size read and write FIFOs and an interface to an external configuration memory. The external configuration memory is used to store parameters specific to the devices to be auto-detected by the core.



SPI-MEM-CTRL

Quad SPI Flash Memory Controller

The SPI-MEM-CTRL core is a highly flexible Quad SPI mode controller and can be configured before synthesis or programmed during run-time to support a large number of SPI Serial Flash memories, even less standard ones. Support for newer serial flash devices with densities larger than 128Mbits is also included.





Alma Technologies S.A.

Leoforos Marathonos 2,
190 09 Pikermi, Greece

T: +30 210 6039 850
F: +30 210 6036 034

info@alma-tech.com
www.alma-tech.com