

CAST



DCT

2-D Forward Discrete Cosine Transform Core

The DCT core implements the 2D Forward Cosine Transform. Most of the image/video compression standards (JPEG, MPEGx, H.261, H.263, DV etc) are based on the Discrete Cosine Transform (DCT). The DCT core, able to operate over 8x8 and 16x16 blocks of samples, covers the needs of hardware image/video compression systems in the most efficient manner. Possibly the fastest core in the market, it is able to provide processing rates up to 190 MSamples/sec in FPGA technologies and over 250 MSamples/sec in ASIC technologies. Furthermore, the core allows the designers to perform area/quality trade-offs by adjusting the cosine coefficients and data-path precision. Finally the 2-4-8 DCT transform, as this is specified in the DVC (or DV) standard, can be optionally supported by the DCT core.

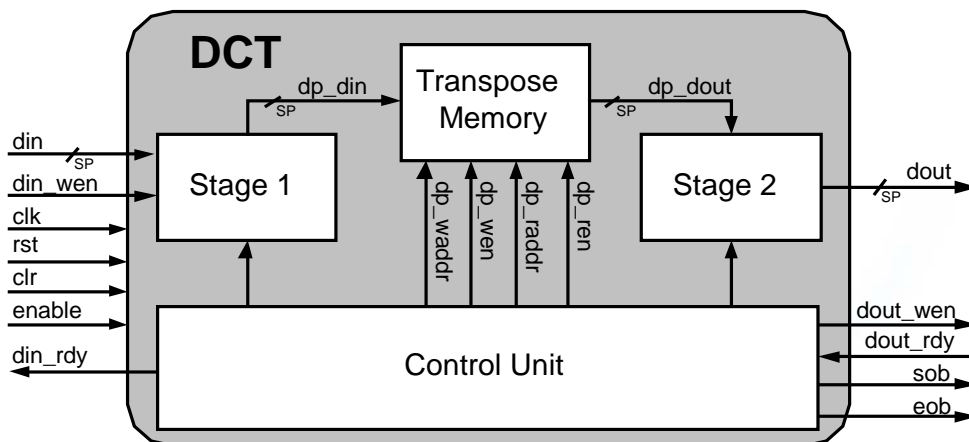
Comprehensive documentation and a complete verification environment - including a bit-accurate model - help designers integrate and verify the core. The DCT is designed for reuse in ASIC and FPGA implementations. The design is fully synchronous with positive edge clocking and no internal tri-state buffers.

Applications

The DCT core can be utilized for a variety of multimedia applications including:

- Office automation equipment (Multifunction printers, digital copiers etc)
- Digital cameras & camcorders
- Video production, video conference
- Surveillance systems

Block Diagram



Features

Ease of Integration & Performance

- High clock speed (>250 MHz in 0.18um ASIC technologies)
- Low gate count
- Single clock cycle per sample operation
- Low latency (87 cycles)

Design Quality

- Fully compliant with the JPEG standard
- Registered input and outputs
- Strictly positive edge triggered fully synchronous design
- Robust verification environment
- No internal latches or tri-states, scan-ready design

Optional add-on Features

- Operation over 16x16 blocks of samples
- Programmable mode of operation (8-8 or 2-4-8)

Functional Description

The forward DCT (DCT) is a transform that converts a signal into its constituent frequency components as represented by a set of coefficients. The inverse DCT (IDCT) reconstructs the original signal from its constituent DCT coefficients. A 2-dimensional array of coefficients results by applying the DCT to 2-dimensional signals, such as images. The core receives image samples and outputs DCT coefficients on a block by block basis, where each block has a size of either 8x8 or 16x16. The core implements the DCT over the input blocks by performing two 1-dimensional transforms, using row-column decomposition, as defined by the following formula:

DCT:

$$Y_{uv} = \frac{2}{N} C_u C_v \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} X_{ij} \cos \frac{(2i+1)u\pi}{2N} \cos \frac{(2j+1)v\pi}{2N} = \sqrt{\frac{2}{N}} C_u \sum_{i=0}^{N-1} \left[\sqrt{\frac{2}{N}} C_v \sum_{j=0}^{N-1} X_{ij} \cos \frac{(2i+1)v\pi}{2N} \right] \cos \frac{(2j+1)u\pi}{2N}$$

where $C_u = C_v = \frac{1}{\sqrt{2}}$ for $u, v = 0$ and $C_u = C_v = 1$ otherwise, X_{ij} are the image samples, Y_{uv} are the DCT coefficients.

The intermediate results being produced from the first 1-dimensional transform are stored in the "Transpose Memory". The Transpose Memory is a dual ported RAM capable of storing an entire 8x8 or 16x16 block resulting from applying the first stage of row decomposition. While the Transpose Memory is written in row-major order, the second stage of processing reads data from the Transpose Memory in a column-major order, effectively performing a transposition of the intermediate results.

The number of bits used for each intermediate result stored in the Transpose Memory, as well as the number of bits used to represent each of the cosine coefficients, is configurable at synthesis time. This allows the designers to perform their own accuracy versus core area tradeoffs. Furthermore, the bit-width of both input image samples and output DCT coefficients is also configurable at synthesis time. It is noted that the default settings for these synthesis parameters, result to a DCT implementation that satisfy the accuracy criteria of the JPEG standard.

The first DCT coefficient of an output block will appear at the output 87 clock cycles after the first image sample of an input block has been fed to the core.

Implementation Results

Reference designs have been evaluated in a variety of technologies. The following are sample Xilinx results obtained after area optimization during synthesis and place and route, while assuming that all core I/Os are routed off-chip.

Xilinx Device	Slices	Frequency	Special Features
Virtex-E v200E-8	1,410	87 MHz	1 BRAM
Spartan2E 2S200-7	1,536	80 MHz	1 BRAM
Spartan3 3S200E-5	960	118 MHz	1 BRAM, 8 Mults
Virtex II 2V250-6	768	154 MHz	1 BRAM, 8 Mults
Virtex II-Pro 2VP2-7	738	189 MHz	1 BRAM, 8 Mults
Virtex-4 4VFX12-11	956	166 MHz	1 BRAM, 8 Mults

Support

The core as delivered is warranted against defects for ninety days from purchase. Thirty days of phone and email technical support are included, starting with the first interaction. Additional maintenance and support options are available.

Verification

The core has been verified through extensive simulation and rigorous code coverage measurements. Being embedded in numerous of products, the core is silicon proven in both FPGA and ASIC technologies.

Deliverables

The core is available in ASIC (synthesizable HDL) and FPGA (netlist) forms, and includes everything required for successful implementation. The Xilinx version includes:

- Post-synthesis EDIF netlist
- A bit-accurate model (BAM) of the core including support of custom test vector generation
- Sophisticated self-checking Testbench (Verilog versions use Verilog 2001) supporting test vectors, expected results, and verification
- RTL and gate level (FPGAs) simulation scripts
- Place and route script
- Comprehensive user documentation, including detailed specifications and a system integration guide