

Ultra-low-power 2D DCT/IDCT architecture using hybrid GDI-PTL adders and Wallace tree multipliers

Bhavya K B¹, A M Vijaya Prakash²

¹Department of Electronics and Communication Engineering, Bangalore Institute of Technology, Visvesvaraya Technological University, Bengaluru, India

²Department of Electronics Engineering, Bangalore Institute of Technology, Visvesvaraya Technological University, Bengaluru, India

Article Info

Article history:

Received May 27, 2026

Revised Aug 23, 2026

Accepted Sep 4, 2026

Keywords:

2D DCT/IDCT

Hybrid gate diffusion input and pass transistor logic

Peak signal-to-noise ratio

Structural similarity index measure

Ultra-low-power area optimization

Wallace tree multiplier

ABSTRACT

The discrete cosine transform (DCT) and inverse discrete cosine transform (IDCT) are widely used in image and video compression, but their hardware implementation requires substantial arithmetic activity and can increase power and area. This work addresses this issue by developing an ultra-low-power, full-precision 2D DCT/IDCT architecture that combines hybrid gate diffusion input and pass transistor logic (GDI-PTL) full adders with Wallace tree multipliers. The objective is to reduce arithmetic power and area while preserving reconstruction fidelity. The proposed architecture uses separable row- and column-wise processing, fixed-point coefficient representation, pipelining, and fine-grained clock and data gating. The register-transfer level (RTL) design is verified using Verilog and MATLAB, and the application specific integrated circuit (ASIC) implementation is synthesized in Cadence Genus and physically implemented in Cadence Innovus Stylus at 45 nm. Post-layout results show a core area of 58,540.8 μm^2 , total power of 2.874 mW at 100 MHz, and a critical data-path delay of 0.261 ns. Image reconstruction is evaluated using compression ratio (CR), peak signal-to-noise ratio (PSNR), and structural similarity index measure (SSIM), demonstrating that the architecture maintains high reconstruction quality while reducing hardware power consumption.

This is an open access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.



Corresponding Author:

Bhavya K B

Department of Electronics and Communication Engineering

Bangalore Institute of Technology, Visvesvaraya Technological University

K.R. road, V.V. Puram, Bengaluru, India

Email: bhavyakb@bit-bangalore.edu.in

1. INTRODUCTION

Growing volumes of multimedia data, real-time video streams, digital images, and medical scans have placed considerable pressure on storage, transmission, and data management infrastructure. Compression techniques, which reduce storage and bandwidth requirements, have therefore become indispensable [1], [2]. Due to excellent energy compaction properties of the discrete cosine transform (DCT), which concentrates most of the signal information into a few low-frequency coefficients, DCT has become the most widely used transform for image compression standards [1], [3]–[6]. The inverse discrete cosine transform (IDCT) is the reverse of this process. It reconstructs the original image data from those coefficients during decompression.

The 2D DCT is computationally intensive, demanding a large number of multiplications and additions, which makes it a natural target for hardware optimization. Traditional DCT/IDCT designs lean heavily on standard complementary metal-oxide-semiconductor (CMOS) arithmetic units, which tend to dissipate significant power and consume considerable chip area as clock frequencies increase. Trimming this

overhead in the arithmetic core is therefore essential to realizing low-power, high-performance hardware for embedded and portable multimedia devices.

In this work, a hybrid gate diffusion input and pass transistor logic (GDI-PTL) full adder-based Wallace tree multiplier is embedded into the 2D DCT/IDCT architecture. Combining GDI and PTL logic styles, the hybrid adder brings down transistor count, lowers switching activity, and reduces dynamic power. Pipelining further boosts throughput, while clock gating and data gating remove power to IDLE functional blocks. The complete system is written in Verilog HDL and synthesized targeting field-programmable gate array (FPGA)/application specific integrated circuit (ASIC) platforms. To keep the implementation manageable, the 2D transform is derived from separable 1D transforms. This is done with the two-stage processing strategy firstly row-wise computation of DCT/IDCT, followed by column-wise transformation. The architecture is aimed at a practical trade-off between speed, power efficiency, and hardware utilization and is a good candidate for real-time image compression and reconstruction.

2. RELATED WORK

The DCT has been adopted as the core operation of image and video compression standards such as JPEG, MPEG, and H.264 for many years, mainly due to its energy compaction property [1]. Despite the widespread use of the standard DCT hardware, it has a significant computational cost and power overhead, which has stimulated a continuous stream of research on lighter, more power-aware architectures. The first attempts in this direction were based on simplifications at algorithm level and approximate transform formulations to decrease the number of arithmetic operations. Brahimi *et al.* [2], Cintra and Bayer [3], Bouguezel *et al.* [4], and Brahimi *et al.* [5] have proposed multiplier-free or approximate 8-point DCT structures that reduce the number of additions and shifts considerably. The trade-off, however, is that the approximation errors accumulate and degrade the quality of the reconstruction, thus limiting their application in scenarios where high fidelity IDCT output is required.

Many research groups have worked on further improving the approximate DCT algorithms in order to reduce the arithmetic workload. Studies [6]–[9] have reported FPGA and ASIC implementations of reduced-addition DCT structures with enhanced throughput and smaller hardware footprints. These designs are typically limited in that they use standard CMOS arithmetic blocks and do not explore much in the way of power savings at the transistor level. Research by Chen *et al.* [10] took a different route with sporadic logarithmic shifters to cut area and power, although the approach grows more complex at higher clock rates and does not take full advantage of lower-power logic families. A related algorithmic direction, explored by Coelho *et al.* [11], reformulates the 2D 8×8 DCT computation using Loeffler factorization and algebraic integer encoding to reduce arithmetic complexity; however, like other algorithm-level techniques, it does not address power reduction at the transistor or circuit level. Improvements have been paid attention to both at the architectural level and at the algorithm level. Vaithyanathan and Seshasayanan [12] proposed a low-power DCT architecture for image compression that improves area and power efficiency through data-flow restructuring and arithmetic reuse. To increase the operating speed, [13]–[15] employed Wallace tree multipliers with carry-select adders. Collectively, these contributions emphasize that the choice of multiplier architecture has a disproportionately large impact on the DCT throughput. However, a common trend of the existing designs is that speed or area tends to be the major target while little attention is paid to fine-grained power management through clock or data gating.

Many efforts have been done at the circuit level to design full adders that can provide low power and high speed, because these cells are the basic building blocks of every multiplier. Techniques such as PTL, GDI, hybrid CMOS–PTL, and XNOR/XOR-based circuits have been reported as effective means of cutting transistor count and switching activity [15]–[22]. Studies [16], [17], [22], [23] showed that 10–12 transistor adder cells can outperform conventional CMOS counterparts in terms of power-delay product (PDP). From [19], [24], [25] similarly reported the practical gains of GDI and hybrid logic in multiplier contexts. A common limitation across this body of work, however, is that evaluations stop at standalone arithmetic units rather than carrying the integration through to a full DCT/IDCT signal processing pipeline. Approximate adders and multipliers have also been examined as a route to ultra-low-power image processing [26], [27]. The energy savings can be substantial, but the numerical errors introduced may be unacceptable in precision-sensitive domains such as medical imaging or high-quality video. Learning-based compression has separately emerged for remote sensing and complex imagery tasks [28], though the computational and memory requirements of such methods make them impractical for real-time, resource-constrained hardware.

In summary, the bulk of existing literature sits in one of two camps: algorithm-level complexity reduction via approximate DCTs, or circuit-level optimization of arithmetic units considered in isolation. Little attention has been paid to jointly designing full-precision DCT architectures with transistor-level low-power logic alongside system-level power management strategies. This work targets precisely that gap, bringing

together a hybrid GDI–PTL full adder–based Wallace tree multiplier within a full-precision 2D DCT/IDCT architecture, complemented by fine-grained clock and data gating. Unlike approximation-based approaches, this design retains full numerical accuracy while delivering meaningful reductions in both dynamic and leakage power, a combination that distinguishes it from prior contributions that address either algorithmic or circuit-level concerns individually, and positions it as a practical candidate for energy-efficient image and video processing. Brahimi *et al.* [5] more recently introduced an 8-point DCT approximation aimed at balancing algorithmic simplicity with acceptable reconstruction quality, increasing compression efficiency in the process. Like the earlier approximate methods, however, the focus remains at the algorithm level rather than the circuit level, so the potential gains from advanced low-power logic styles go unrealized. Separately, at the transistor level, Naseri and Timarchi [22] developed optimized XOR and XNOR gate structures to produce a full adder with notably better PDP than standard CMOS cells. Their work confirms the value of transistor-level arithmetic optimization but stops short of integrating these cells into a larger processor such as a DCT/IDCT system.

Unlike the directions taken in existing literature, this work attempts to close the gap between algorithm-level efficiency and transistor-level power optimization by embedding a hybrid GDI–PTL arithmetic design within a full-precision 2D DCT/IDCT architecture. Whereas approximation-based DCT methods [18] trade reconstruction accuracy for power savings, the design here keeps full numerical precision and instead reduces power through lower transistor count and reduced switching activity. Block-level clock gating then works alongside the low-power arithmetic to yield further reductions in both dynamic and leakage components. Instead of having to select either algorithmic simplicity or circuit optimization, the architecture aims at achieving both simultaneously—a unique feature compared to existing solutions which makes it ideal for processing images or videos with limited power resources.

3. METHOD

The proposed block diagram in Figure 1 gives an overall picture of the low-power 2D DCT–IDCT system developed for image compression and reconstruction. Raw image data enter the system already partitioned into fixed-size blocks; each block is then pushed through a row–column separable 2D DCT, shifting the representation from the spatial to the frequency domain. The heavy lifting, multiply–accumulate operations that DCT requires is done by a Wallace tree multiplier whose full adder cells are realized in hybrid GDI–PTL logic.




Figure 1. Block diagram of the proposed low-power 2D DCT–IDCT architecture with hybrid GDI–PTL Wallace tree multipliers, clock gating, data gating, and MATLAB-based verification

This approach directly cuts transistor count and reduces switching activity, both of which contribute to lower power figures. Pipeline registers ensure that data flow at the desired clock speed without the critical path being extended during the intervals between the computation phases. For controlling the rest of the dynamic power consumption, clock gating, and data gating techniques are used throughout all the DCT and IDCT stages to avoid spurious activity when the pipeline is not in use or invalid data is present. As soon as the DCT computations are done, its coefficient results are fed directly into the 2D IDCT stage. In order to avoid any duplication of hardware, IDCT operations are carried out with the help of the same arithmetic units by transposing the coefficients, an approach that has been adopted for reducing both the power and area costs.

At last, the decoded image is compared with the MATLAB output image with the help of peak signal-to-noise ratio (PSNR) and compression ratio (CR).

3.1. Overview of the proposed 2D discrete cosine transform/inverse discrete cosine transform architecture

This approach leads to an energy-efficient implementation of the 2D DCT/IDCT, where each transform operates on a standard 8×8 block of image pixels with 64 data samples at a time. In order to increase the throughput and minimize chip area and power consumption of the design, we have exploited the fact that the DCT is separable, thus splitting one large and complex 2D transform into two simpler transforms along rows and columns. Processing first rows and next columns makes the solution practically feasible, as computations become less complex while silicon area can be minimized. The whole system is implemented using a parameterized Verilog module (gdi_2d_dct_idct). Input pixel width, coefficient precision, fractional scaling, and output resolution are all parameterized in this module so that this register-transfer level (RTL) core can be used repeatedly across different technology nodes just by modifying some parameters to fit certain application needs.

3.2. Fixed-point discrete cosine transforms formulation and coefficient modeling

To ensure optimal trade-off between hardware efficiency and precision, the fixed-point form of the DCT algorithm has been selected. The cosine coefficients are stored in Q format (Q14) with a coefficient size of 16 bits. The coefficient function (dctA) is predefined to compute orthonormal coefficients of DCT from frequency index and spatial index combination. For inverse DCT operations, the same set of coefficients are used with the help of transposition technique (dctM). Using the same set of coefficient tables helps to minimize space and power without compromising performance.

3.3. Hybrid gate diffusion input and pass transistor logic full adder design

The key arithmetic component in this system architecture design is the GDI-PTL full adder. This type of full adder is employed in an attempt to realize power and propagation delay reductions simultaneously. GDI cuts down the power dissipated by ensuring fewer transistors while PTL ensures fast signal transmission with minimal capacitance requirements. With around 10 transistors, versus the 28 needed by a standard CMOS full adder, the area and power savings are substantial. Although PTL logic carries a reduced voltage swing, output restoration is applied to maintain signal integrity through cascaded stages. In pipelined arithmetic datapaths these cells also exhibit lower switching activity, which keeps dynamic power in check. The schematic of the resulting hybrid GDI-PTL full adder cell, adapted from our earlier work [23], is shown in Figure 2.




Figure 2. Schematic diagram of the hybrid GDI-PTL full adder cell [23]

3.4. Hybrid Wallace tree multiplier architecture

In a DCT/IDCT pipeline, multiply-accumulate operations between pixel samples and cosine coefficients account for the bulk of the workload, which makes the multiplier the single most important block for meeting performance targets. To carry this load, the architecture uses a hybrid Wallace tree multiplier, assembled from the hybrid GDI-PTL full-adder cells described earlier. All partial products are generated in parallel and then fed into a Wallace tree reduction network that compresses them down to two operands in as few levels as possible; this is what keeps the adder tree shallow and cuts the critical path well below what a

conventional array multiplier would need. The operand width is configurable, so the same multiplier block handles pixel-coefficient and intermediate-coefficient multiplications without any structural change. Pipelining across the multiplier stages means higher clock rates are achievable without a proportional rise in power. The net result of combining Wallace tree parallelism with hybrid adder efficiency is a multiplier that sustains high throughput with lower dynamic and leakage power.

3.5. Two-stage 2D discrete cosine transform/inverse discrete cosine transform computation

3.5.1. Input unpacking and row transformation

The 64-element input vector is expanded into an internal linear array that represents an 8×8-pixel matrix. In the row transformation stage, each row is handled in turn. Pixel values for a given row are multiplied against the appropriate DCT coefficients through the hybrid Wallace tree multipliers, then a balanced adder tree accumulates the partial products. The accumulated values are scaled using fixed-point scaling and arithmetic right shift operations to handle fractional coefficient scaling. The output values from the rows are stored in internal registers (Y_flat) using enough bits to avoid overflow.

3.5.2. Column transformation and output generation

In the second stage, the intermediate row results are processed column-wise using the same multiplier array. For IDCT operation, transposed coefficients are automatically selected, enabling full hardware reuse. The column-stage outputs undergo rounding and saturation logic to constrain results within the desired output bit width, ensuring numerical stability and preventing overflow. The final 64 DCT/IDCT coefficients or reconstructed pixel values are packed into a single output vector (data_out) for efficient downstream processing.

3.6. Finite state machine-based control and pipelining strategy

A centralized finite state machine (FSM) governs the overall operation of the architecture, transitioning through four states: IDLE, ROW_DCT, COL_DCT, and DONE. The FSM ensures correct sequencing of data flow while synchronizing computation and register updates. The design completes one 2D DCT/IDCT operation in a deterministic number of clock cycles, with row and column computations executed in successive stages. Pipelining between these stages allows partial overlap of computation, improving throughput without increasing control complexity. The resulting FSM-controlled datapath architecture is illustrated in Figure 3.




Figure 3. FSM-controlled architecture of the proposed low-power 2D DCT/IDCT module using hybrid Wallace multipliers

3.7. Clock gating and low-power optimization

Keeping dynamic power low calls for careful management of clock activity, so the design uses block-level fine-grained clock gating throughout the pipeline. Clock-enable signals (en_row and en_col) are driven directly by FSM state and gate the row-wise and column-wise arithmetic units so they only receive a clock when actively computing. During IDLE and DONE states, every arithmetic block is held off, eliminating spurious clock transitions. Alongside clock gating, data gating is applied to input and intermediate registers to block needless data propagation. Both techniques reduce switching activity, which is the primary power cost in arithmetic-heavy designs. Pairing them with the low transistor count of the hybrid GDI-PTL arithmetic units produces complementary reductions across the design, bringing both dynamic and leakage power down together—a property that makes the architecture well suited to power-sensitive image and video processing deployments.

3.8. Verification and physical design flow

The architecture is functionally verified using Verilog testbenches in ModelSim and Vivado environments. Verification includes correctness validation for both DCT and IDCT operations and visual quality assessment of reconstructed images. For ASIC implementation, the synthesized netlist is generated using Cadence Genus, followed by floor planning, placement, clock tree synthesis (CTS), routing, and signoff using Cadence Innovus. Post-layout extraction enables accurate evaluation of power, delay, and area, validating the effectiveness of the proposed hybrid logic and clock-gating methodology.

4. RESULTS AND DISCUSSION

The 2D DCT/IDCT architecture was evaluated across three levels: algorithmic validation in MATLAB, RTL-level functional verification, and post-layout power and timing analysis in Cadence Innovus Stylus. The MATLAB program was employed initially to check the validity of DCT and IDCT algorithms. Subsequently, PSNR, structural similarity index measure (SSIM), and CR techniques were used to assess the reconstructed images and the results proved to be free from reconstruction error, an accuracy that is indispensable in the context of medical applications. RTL level implementation using Verilog generated the DCT coefficients as well as the reconstructed image after IDCT process, which matched closely to those of the MATLAB software results. Minor differences from software and hardware implementations confirm good fixed-point representation, proper rounding and saturation mechanisms, as well as correct order of state machine operations in row and column transformations.

Power values post-layout obtained using Cadence Innovus software demonstrate that hybrid GDI-PTL Wallace tree multiplier significantly cuts down on the number of switches in comparison with standard arithmetic circuits. The technique of clock gating eliminates unnecessary clock signals when they are not required, while data gating ensures that there are no redundant transitions performed in the input registers and in-between registers. A comparison between these power values and MATLAB image quality measures proves that the power savings do not affect reconstruction accuracy. Thus, the proposed architecture ensures achieving all critical performance parameters simultaneously.

Figures 4–6 contain the images that were utilized in the experiment. Figure 4 contains some common benchmark test images that can be found in literature concerning image processing. These images have been carefully selected in order to incorporate a variety of scene properties such as edges, textures, smoothness, and contrasts. The images in Figure 5 are sample color images from the USC-SIPI image database in different spatial dimensions. The corresponding gray-scale images from the USC-SIPI image database in different resolutions are shown in Figure 6.




Figure 4. Sample standard test images




Figure 5. Sample color images for USC-SIPI dataset of different dimensions




Figure 6. Sample gray-scale images for USC-SIPI dataset of different dimensions

Image reconstruction using the MATLAB based IDCT method is compared with the image reconstruction using the proposed Verilog IDCT design along with the actual input image, as shown in Figure 7. The image reconstruction using MATLAB acts as the software version that helps to show near-perfect image reconstruction with negligible distortion in the output image. The output image from the Verilog IDCT image reconstruction is similar to the one obtained from MATLAB image reconstruction with important features like edge, texture, and luminance intact.




Figure 7. Visual comparison of IDCT image reconstruction using MATLAB and proposed Verilog architecture

Table 1 presents the compression results for different standard test images. Images with smoother or more correlated content achieve higher CRs and better reconstruction quality, while images with edges, textures, or low contrast show lower compression efficiency. Despite this variation, all images maintain acceptable PSNR and SSIM values, indicating good reconstruction quality across different image types.

Table 1. Compression results for standard data base images

Image	Image type	Image dimension	Coefficient-based CR	Bit-accurate CR	PSNR (dB)	SSIM
Lena.jpg	Natural	256×256	32.00:1	16.00:1	38.48	0.973
cameraman.tif	Edges	256×256	1.02:1	0.51:1	30.4	0.9478
peppers.png	Smooth	256×256	1.00:1	0.50:1	33.02	0.9316
coins.png	Shadows	512×512	1.01:1	0.50:1	31.67	0.943
rice.png	Low contrast	512×512	1.00:1	0.50:1	32.58	0.9259
mri.tif	Medical	512×512	2.53:1	1.27:1	38.98	0.9732
saturn.png	Texture	512×512	1.15:1	0.58:1	35.59	0.9708

This lossless reconstruction behavior is illustrated in Figure 8, which compares the original image against the images reconstructed using MATLAB and the proposed hardware IDCT implementation. The figure demonstrates the visual similarity between the original image and the reconstructed images obtained using MATLAB and the proposed Verilog-based IDCT architecture. The negligible perceptual difference confirms accurate hardware implementation and validates the lossless or near-lossless reconstruction capability of the proposed design.




Figure 8. Lossless image reconstruction validation using MATLAB and hardware IDCT implementation

Figures 9 and 10 show reconstruction results using medical image databases. In each panel of Figure 9, Figure 9(a) shows the original MRI slice from the Kaggle brain-tumor dataset, while Figure 9(b) shows the corresponding image reconstructed by the proposed hardware IDCT path, with its computed PSNR: 27.78 dB for Image 3, 36.45 dB for Image 10, 36.01 dB for Image 16, 36.44 dB for Image 45, 34.05 dB for Image 68, and 34.26 dB for Image 70. Similarly, in Figure 10, Figure 10(a) shows the original retinal fundus image from the DRIVE dataset and Figure 10(b) shows the corresponding hardware-reconstructed image, with PSNR values of 31.52 dB, 31.38 dB, 33.02 dB, 32.83 dB, 34.57 dB, and 33.61 dB for the 01_test, 02_test, 08_test, 13_test, 15_test, and 20_test images, respectively. In both figures, one can see that output images from hardware follow the originals well, preserving the structural boundaries and intensity gradient levels that are clinically important for tumor margin and vessel-map interpretation. The quantitative measurement of PSNR values listed near the corresponding images further confirms our claims about successful performance of Verilog DCT/IDCT based architecture in terms of reconstruction quality for images representing various medical imaging types. Because diagnostic decisions in MRI and retinal imaging depend on fine structural detail, PSNR and SSIM values in this range indicate that the compression-reconstruction pipeline does not introduce artifacts likely to obscure clinically relevant features. Therefore, the proposed hardware architecture may be considered as a strong candidate for image reconstruction tasks in medical imaging.




Figure 9. Reconstruction results for six representative images from the Kaggle brain-tumor-classification-MRI dataset, comparing; (a) the original image and (b) the corresponding uncompressed/reconstructed image with its PSNR value




Figure 10. Reconstruction results for six representative retinal fundus images from the DRIVE dataset, comparing; (a) the original image and (b) the corresponding uncompressed/reconstructed image with its PSNR value

Table 2 shows a comprehensive performance evaluation for DCT-based image compression architecture when applied to medical images from Kaggle MRI glioma tumor database, USC-SIPI, and DRIVE retinal image databases. It presents the coefficients-based CR, bit-accurate CR, space saving percentage, and reconstructed image quality (PSNR and SSIM). According to the table, space savings can reach up to 62.12%, while providing high reconstructed image quality (above 30 dB PSNR value and above 0.95 SSIM for the majority of tested images).

The CR based on the coefficient is obtained through taking the quotient of the total number of pixels divided by the number of non-zero DCT coefficients. The bit-accurate CR is more comprehensive because it takes into consideration not only the non-zero DCT coefficients but also the word-length required to represent such coefficients. Savings in space are computed based on the CR using coefficients. What should be emphasized is the fact that the compression algorithm takes precedence over the need for good performance in terms of reconstruction accuracy and simplicity in design rather than the aggressiveness in compression. Hence, in some cases, the bit-accurate CR may be less than one.

Table 2. Compression performance metrics for different image databases

Database	Image	Image dimension	Coefficient-based CR	Bit-accurate CR	Space saving (%)	PSNR	SSIM
Kaggle MRI glioma_tumor	Image(3)	256×256	2.64:1	0.82:1	62.12	37.78	0.9934
	Image(6)	256×256	1.64:1	0.72:1	39.02	36.78	0.9134
	Image(63)	512×512	1.62:1	0.81:1	38.27	33.46	0.9515
	Image(68)	512×512	1.55:1	0.77:1	35.48	34.05	0.9517
	Image(70)	512×512	1.56:1	0.78:1	35.9	34.26	0.952
	Image(19)	225×225	1.05:1	0.15:1	4.76	33.31	0.994
	Image(16)	398×397	1.16:1	0.29:1	13.79	36.01	0.9863
USC-SIPI	Image(45)	384×382	2.40 :1	0.60:1	58.33	36.44	0.9926
	<u>4.1.05</u>	256×256	1.76:1	0.61:1	43.18	32.55	0.8634
	<u>4.1.07</u>	256×256	2.50:1	0.72:1	60	28.67	0.4562
	<u>4.2.03</u>	512×512	1.93:1	0.73:1	48.19	27.62	0.5936
	<u>4.2.07</u>	512×512	1.08:1	0.45:1	7.41	38.52	0.7534
	4.02.05	512×512	2.11:1	0.87:1	52.61	37.57	0.9653
	7.1.01	512×512	1.22:1	0.54:1	18.03	39.49	0.97695
	<u>boat.512</u>	512×512	1.70:1	0.58:1	41.18	35.75	0.97403
	5.3.01	1024×1024	2.31:1	0.57:1	56.71	37.33	0.97769
	5.3.02	1024×1024	1.84:1	0.82:1	45.65	35.18	0.96739
	<u>5.1.12</u>	256×256	1.30:1	0.56:1	23.08	35.06	0.9861
	<u>5.1.10</u>	256×256	2.23:1	0.92:1	55.16	31.65	0.97162
DRIVE (Retina)	01_test.tif	565×584	1.03:1	0.53:1	2.91	31.52	0.9941
	02_test.tif	565×584	1.09:1	0.69:1	8.26	31.38	0.9723
	08_test.tif	565×584	1.13:1	0.63:1	11.5	33.02	0.9725
	13_test.tif	565×584	1.14:1	0.51:1	12.28	32.83	0.9716
	15_test.tif	565×584	1.11:1	0.71:1	9.91	34.57	0.9728
	20_test.tif	565×584	1.10:1	0.80:1	9.09	33.61	0.9721

4.1. Performance evaluation metrics for discrete cosine transform/inverse discrete cosine transform architectures

To quantitatively evaluate the effectiveness of the proposed image compression architecture, three standard metrics are employed: CR, PSNR, and SSIM. These metrics collectively capture compression efficiency, numerical reconstruction fidelity, and perceptual image quality.

- a. CR: CR quantifies the data reduction achieved by transform-based encoding and is defined as:

$$CR = \frac{\text{Total bits in original image}}{\text{Total bits after compression}}$$

For a fair comparison, CR is evaluated identically for DCT/IDCT architecture using the quantization and coefficient selection strategy. Two forms of CR are reported:

- Coefficient-based CR: this metric considers only the number of retained DCT coefficients after quantization.

$$CR_{\text{Coeff}} = \frac{N_{\text{total coeff}}}{N_{\text{non-zero coeff}}}$$

- Bit-accurate CR: this accounts for the actual number of bits required to represent the compressed data, including coefficient magnitudes and sign bits:

$$CR_{bit} = \frac{\text{Bits in original image}}{\text{Bits used for compressed coefficients}}$$

Since both DCT-only and DCT/IDCT architectures use the same forward DCT and quantization stages, the CR remains identical for both methods, ensuring a structurally fair comparison.

- b. PSNR: PSNR measures the reconstruction accuracy by comparing the reconstructed image with the original image. It is derived from the mean squared error (MSE):

$$MSE = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N [I(i,j) - I'(i,j)]^2$$

where $I(i,j)$ is the original image and $I'(i,j)$ is the reconstructed image. For 8-bit grayscale images, PSNR is computed as:

$$PSNR(dB) = 10 \log_{10} \left(\frac{255^2}{MSE} \right)$$

In the DCT/IDCT architecture reconstructs the image using an approximate or truncated inverse operation, leading to a controlled reduction in PSNR. This trade-off is intentional and enables significant reductions in hardware complexity and power consumption.

- c. SSIM: SSIM evaluates perceptual image quality by comparing structural information between the original and reconstructed images. SSIM values range from 0 to 1, with values closer to 1 indicating better perceptual similarity.

The FSM-driven sequencing of control and data signals during a representative computation cycle is captured in the simulation waveform of Figure 11. Functional verification of the proposed 2D DCT/IDCT architecture was carried out through RTL simulation using database image pixel values applied as a hexadecimal input file. The simulation waveform confirms correct sequencing of control and data signals under the supervision of the FSM. As observed, the input unpacking, row-wise and column-wise DCT/IDCT computations, and output packing occur in a synchronized manner, with valid output samples generated only after completion of the required computation cycles. The absence of data hazards and timing violations in the waveform demonstrates stable pipeline operation and correct handling of block-based image data. These results validate the functional correctness of the proposed architecture prior to synthesis and physical implementation. Table 3 benchmarks the post-layout power, area, and PDP of the proposed unified DCT/IDCT design against previously reported DCT-only architectures.




Figure 11. Functional simulation waveform illustrating the operation of the proposed 2D DCT/IDCT architecture for a database image

The layout in Figure 12 shows standard-cell implementation with optimized placement, clock-tree synthesis, and routed interconnects incorporating clock and data gating techniques. The compact and regular structure validates the area efficiency and physical realizability of the proposed architecture. Post-layout implementation of the proposed low-power 2D DCT/IDCT architecture was carried out using Cadence Innovus Stylus following a standard digital physical design flow, including floorplanning, placement, CTS, routing, and sign-off analysis. The finalized layout demonstrates compact area utilization, achieved through efficient reuse of arithmetic resources and the reduced transistor count of hybrid GDI-PTL based Wallace tree

multipliers. The core area reflects the regular structure of the row-column DCT/IDCT datapath and effective standard-cell placement with minimal routing congestion.

Table 3. Comparison of power, area, and PDP for reported DCT-only architectures versus the proposed unified DCT/IDCT design

Work	Technology (nm)	Architecture	Total power (mW)	Freq. (MHz)	Power (mW/MHz)	Area (μm^2)	Data width (bits)	Area ($\mu\text{m}^2/\text{bit}$)	Critical path delay (ns)	PDP (pJ)
Dhandapani and Ramachandran [21]	90	Approximate DCT (adder-based)	0.1954	100	0.00195	1,954	1024	1.91	0.616	0.12
Cintra and Bayer [3]	90	Low-power DCT	0.3173	—	—	2,459	1024	2.4	0.103	0.033
Mendez <i>et al.</i> [29]	45	Improved B.G. Lee DCT	9.72	—	—	17,690	1024	17.28	26.48	257.38
Senthilkumar <i>et al.</i> [30]	90	Computation-reduced DCT	9.2727	—	—	57,829	1024	56.47	—	—
Proposed work (post-layout)	45	Hybrid GDI-PTL + Wallace (Exact DCT/IDCT)	2.874	100	0.0287	58,540.80	1024	57.17	0.261	0.75




Figure 12. Post-place-and-route physical layout of the proposed 2D DCT/IDCT architecture using Cadence Innovus Stylus

Power analysis performed after routing indicates a significant reduction in dynamic power consumption, primarily due to the integration of clock gating and data gating across DCT and IDCT stages. CTS was optimized to minimize clock skew and insertion delay, ensuring balanced clock distribution while reducing unnecessary clock switching. Timing analysis confirms that the design meets setup and hold constraints, with positive worst negative slack (WNS), validating the effectiveness of pipelining in reducing critical path delay. Overall, the post-layout results confirm that the proposed architecture is physically realizable and suitable for low-power VLSI implementation.

Table 4 summarizes the key implementation parameters obtained from pre-layout RTL synthesis and the final post-layout design, highlighting the impact of physical implementation on area, power, and timing. Using RTL synthesis (using Cadence Genus, total cell area is 1,679,792.12 μm^2 with 26,341 standard cells, where combinational logic accounts for approximately 91.1% of the total area). Pre-layout power analysis indicates total power consumption is 97.17 mW (almost all dynamic in nature), demonstrating the significance of the pipelining and clock/data gating techniques employed to minimize switching activity in the design. Post-

layout implementation (using Cadence Innovus Stylus) confirms the physical viability of the design. After placement, CTS, and routing, the finalized design occupies a core area of 58,540.8 μm^2 , consists of 17,221 standard cells, and contains 57,057 logic gates. Timing closure has been reached with a clock period of 10 ns, with a data path delay of only 0.261 ns and a worst-case positive slack value of 7.639 ns, allowing for ample margin for increased operational frequency. Post-layout power measurements indicate total power dissipation of 2.874 mW at supply voltage (1.08 V) with approximately 99.9% of total power being dynamic power. These results support the conclusion that the hybrid GDI-PTL and Wallace tree architectures offer significant power efficiency without compromising timing margin or image reconstruction fidelity.

Table 4. Performance comparison of RTL and post-layout implementations

Parameter	RTL (Genus-pre-layout)	Post-layout (Innovus Stylus)
Technology	45 nm CMOS	45 nm CMOS
Core area (μm^2)	1,679,792.12	58,540.8
Standard cells	26,341	17,221
Total gates	-	57,057
Total power (mW)	97.17	2.874
Leakage power (mW)	0.078	0.00239
Dynamic power (mW)	97.09	2.871
Supply voltage (V)	1.08	1.08
Clock period (ns)	10	10
Worst-case slack (ns)	6.42	7.639
Image quality	Bit-accurate	Bit-accurate

4.2. Limitations, scalability, and deployment considerations

Relative to the reported DCT-only baselines in Table 3, the proposed unified DCT/IDCT design demonstrates a favorable power efficiency at the same 45 nm technology node. When normalized to data width, the 45 nm baseline of Mendez *et al.* [29] dissipates approximately 3.4 times more power per bit than the proposed architecture. The comparison with the 90 nm designs is more varied, with some reported architectures exhibiting lower power and others higher power than the proposed design. These differences are influenced by technology-node scaling, architectural characteristics, and implementation techniques, including the hybrid GDI-PTL adder cells, Wallace tree reduction, and fine-grained clock/data gating employed in the proposed design. Moreover, unlike the approximate or computation-reduced DCT architectures listed in Table 3, the proposed design performs the full, non-approximate 2D DCT and IDCT, providing this power-efficiency advantage without relying on transform approximation or reduced computation.

Embedded and real-time multimedia relevance: the FSM-based control, pipelined datapath, and clock/data-gated arithmetic units make the architecture directly applicable to embedded image and video front ends, such as low-power medical imaging capture modules, portable diagnostic devices, and edge-based multimedia encoders, where both battery life and reconstruction fidelity are constrained. The design's block-based 8×8 processing and deterministic per-block latency are compatible with the block-processing structure of JPEG-, MPEG-, and H.264-style pipelines, allowing it to be integrated as the transform stage of a larger embedded codec without restructuring the surrounding encode/decode flow.

Trade-offs and potential weaknesses: clock and data gating reduce switching power but introduce enable-signal generation and isolation logic that adds a small area and timing overhead, and, if the gating condition is not registered carefully, can introduce glitches on the gated clock or data path. In this design, gating is derived synchronously from FSM state (en_row, en_col), which avoids combinational glitches on the clock enable at the cost of one cycle of gating latency when transitioning between states; this is a deliberate trade-off of a marginal throughput overhead for glitch-free low-power operation. At substantially higher target clock frequencies than the 100 MHz evaluated here, the reduced voltage swing inherent to PTL-based signal transmission would be expected to narrow the pipeline's timing margin before the hybrid adder's power advantage is lost, so operation at very high frequency is a limitation of the present design rather than a demonstrated capability. Likewise, the results reported here are for 8×8 block processing typical of JPEG-style compression; scaling to larger block sizes or to full video frames would require either replicating the datapath or time-multiplexing it across blocks, and the resulting area/power/throughput trade-off has not been characterized in this work.

Scalability to higher resolutions and video streams: because the architecture processes image data in independent 8×8 blocks, extending it to larger still images requires only sequential or parallel replication of the existing block pipeline, with no change to the per-block hybrid GDI-PTL/Wallace tree datapath. Extending the design to video would additionally require a frame buffer and block-scheduling logic to sustain the required frame rate, and the clock/data gating strategy would need to be re-evaluated against the higher and more

continuous switching activity typical of video streams; this is identified as a direction for future work rather than a result demonstrated in the present study.

Deployment considerations: integrating the proposed transform core into an existing JPEG/MPEG/H.264 pipeline is expected to be straightforward at the architectural level, since those standards already assume block-based 2D DCT/IDCT, but a full codec integration (including entropy coding, quantization tables, and rate control) has not been implemented or evaluated here. Fabrication cost and portability across technology nodes were not directly assessed; the RTL is parameterized (subsection 3.1) to ease retargeting, but power, area, and timing figures reported in section 4 are specific to the 45 nm node used for physical implementation and would need to be re-characterized for other nodes.

Future work: natural extensions of this study include FPGA prototyping for real-time throughput validation, extension of the hybrid GDI-PTL/Wallace tree datapath to video codec pipelines, and exploration of a configurable exact/approximate mode that would allow the same hardware to trade a small, bounded amount of reconstruction accuracy for further power reduction in non-diagnostic use cases, while retaining the present full-precision mode for medical imaging. From a standards and policy perspective, medical imaging guidelines that require near-lossless or lossless compression for diagnostic storage and transmission (e.g., DICOM-based archiving) are a natural application context for a design of this kind, since it targets low power without sacrificing the numerical accuracy such guidelines require.

5. CONCLUSION

This study outlines the design and implementation of a two-dimensional DCT/IDCT processor. It starts from validation via algorithm development, RTL level synthesis, and finally to post-layout physical implementation. The accuracy of reconstruction has been proven by high CR, PSNR, and SSIM values using MATLAB testing. The results from Cadence Genus show an effective arithmetic operation for implementation, while post-layout results using Cadence Innovus Stylus prove that there is clean timing closure with ultra-low power consumption. This design occupies a core region of 58,540.8 μm^2 ; it uses only 2.874 mW operating voltage of 1.08 V with a big positive timing slack of 7.639 ns. These findings provide concrete proof of the effectiveness of hybrid GDI-PTL adders, Wallace tree multipliers, and clock and data gating. Thus, consistency in MATLAB, RTL, and physical implementations in all results shows that this architecture could be reliable and efficient for real-time processing.

ACKNOWLEDGMENTS

The authors would like to thank the Department of Electronics and Communication Engineering, Bangalore Institute of Technology for providing the necessary facilities and support to carry out this research work.

FUNDING INFORMATION

Authors state no funding involved.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Bhavya K B	✓	✓	✓	✓	✓	✓		✓	✓	✓			✓	
A M Vijaya Prakash		✓				✓		✓	✓	✓	✓	✓		

C : **C**onceptualization

M : **M**ethodology

So : **S**oftware

Va : **V**alidation

Fo : **F**ormal analysis

I : **I**nvestigation

R : **R**esources

D : **D**ata Curation

O : Writing - **O**riginal Draft

E : Writing - Review & **E**ditng

Vi : **V**isualization

Su : **S**upervision

P : **P**roject administration

Fu : **F**unding acquisition

CONFLICT OF INTEREST STATEMENT

Authors state no conflict of interest.

DATA AVAILABILITY

The image datasets used in this study are publicly available from the respective dataset repositories, including the USC-SIPI Image Database, the Kaggle brain-tumor MRI dataset, and the DRIVE retinal image database. No proprietary dataset was used. The processed results generated during this study are available from the corresponding author upon reasonable request.

REFERENCES

- [1] G. K. Wallace, "The JPEG still picture compression standard," *IEEE Transactions on Consumer Electronics*, vol. 38, no. 1, pp. xviii–xxxiv, 1992, doi: 10.1109/30.125072.
- [2] N. Brahimi, T. Bouden, T. Brahimi, and L. Boubchir, "Lossy image compression based on efficient multiplier-less 8-points DCT," *Multimedia Systems*, vol. 28, no. 1, pp. 171–182, Feb. 2022, doi: 10.1007/s00530-021-00762-0.
- [3] R. J. Cintra and F. M. Bayer, "A DCT Approximation for Image Compression," *IEEE Signal Processing Letters*, vol. 18, no. 10, pp. 579–582, Oct. 2011, doi: 10.1109/LSP.2011.2163394.
- [4] S. Bouguezel, M. O. Ahmad, and M. N. S. Swamy, "A fast 8x8 transform for image compression," in *Proceedings of the International Conference on Microelectronics, ICM*, IEEE, Dec. 2009, pp. 74–77, doi: 10.1109/ICM.2009.5418584.
- [5] N. Brahimi, T. Bouden, T. Brahimi, and L. Boubchir, "A novel and efficient 8-point DCT approximation for image compression," *Multimedia Tools and Applications*, vol. 79, no. 11–12, pp. 7615–7631, Mar. 2020, doi: 10.1007/s11042-019-08325-2.
- [6] U. S. Potluri, A. Madanayake, R. J. Cintra, F. M. Bayer, S. Kulasekera, and A. Edirisuriya, "Improved 8-Point Approximate DCT for Image and Video Compression Requiring Only 14 Additions," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 61, no. 6, pp. 1727–1740, Jun. 2014, doi: 10.1109/TCSI.2013.2295022.
- [7] D. Vaithiyanathan and P. Dheepansundaravelu, "Implementation of 8-Point Approximate DCT for Image Compression," *International Journal of Advanced Research in Computer Science and Software Engineering*, vol. 6, no. 9, pp. 441–448, 2016.
- [8] D. Vaithiyanathan, R. Seshasayanan, S. Anith, and K. Kunaraj, "A low-complexity DCT approximation for image compression with 14 additions only," in *2013 International Conference on Green Computing, Communication and Conservation of Energy (ICGCE)*, IEEE, Dec. 2013, pp. 303–307, doi: 10.1109/ICGCE.2013.6823450.
- [9] J. B. Pari and D. Vaithiyanathan, "An Optimized FPGA Implementation of DCT Architecture for Image and Video Processing Applications," in *2019 International Conference on Wireless Communications Signal Processing and Networking (WiSPNET)*, IEEE, Mar. 2019, pp. 186–191, doi: 10.1109/WiSPNET45539.2019.9032787.
- [10] J. Chen, Y. Wang, J. Zhao, and S. Rahardja, "A new area and power efficient DCT circuits using sporadic logarithmic shifters," *IEICE Electronics Express*, vol. 16, no. 14, p. 20190317, 2019, doi: 10.1587/elex.16.20190317.
- [11] D. F. G. Coelho, S. Nimmalapalli, V. S. Dimitrov, A. Madanayake, R. J. Cintra, and A. Tisserand, "Computation of 2D 8x8 DCT Based on the Loeffler Factorization Using Algebraic Integer Encoding," *IEEE Transactions on Computers*, vol. 67, no. 12, pp. 1692–1702, Dec. 2018, doi: 10.1109/TC.2018.2837755.
- [12] D. Vaithiyanathan and R. Seshasayanan, "Low power DCT architecture for image compression," in *2013 International Conference on Advanced Computing and Communication Systems*, IEEE, Dec. 2013, pp. 1–6, doi: 10.1109/ICACCS.2013.6938745.
- [13] S. Anju and M. Saravanan, "High Performance DCT Implementation Using Reduced Complexity Wallace Multiplier and High Speed Carry Select Adder," *International Journal of Advanced Research in Computer and Communication Engineering*, vol. 2, no. 4, pp. 1809–1813, 2013.
- [14] T. Sribhasavi, K. Radhika, and G. Karthick, "Design and Analysis of Low Power High Speed Combinational Multiplier and Wallace Tree Multiplier for Image Compression," *International Journal of Advanced Science and Technology*, vol. 29, no. 3, pp. 3675–3684, 2020.
- [15] Y. d. Ykuntam, K. Pavani, and K. Saladi, "Design and analysis of High speed wallace tree multiplier using parallel prefix adders for VLSI circuit designs," in *2020 11th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, IEEE, Jul. 2020, pp. 1–6, doi: 10.1109/ICCCNT49239.2020.9225404.
- [16] M. Ahmad, K. Manjunathachari, and K. Lalkishore, "Design and analysis of low run-time leakage in a 10 Transistors full adder in 45nm technology," in *2016 IEEE Region 10 Conference (TENCON)*, IEEE, Nov. 2016, pp. 152–156, doi: 10.1109/TENCON.2016.7847979.
- [17] J.-F. Lin, Y.-T. Hwang, and M.-H. Sheu, "Low power 10-transistor full adder design based on degenerate pass transistor logic," in *2012 IEEE International Symposium on Circuits and Systems*, IEEE, May 2012, pp. 496–499, doi: 10.1109/ISCAS.2012.6272074.
- [18] C. P. Kadu and M. Sharma, "Area-Improved High-Speed Hybrid 1-bit Full Adder Circuit Using 3T-XNOR Gate," in *2017 International Conference on Computing, Communication, Control and Automation (ICCUBEA)*, IEEE, Aug. 2017, pp. 1–5, doi: 10.1109/ICCUBEA.2017.8463827.
- [19] K. Murugan, R. Nithya, K. Prasanth, S. Fowjiya, R. U. Mageswari, and E. A. M. Ali, "Analysis of Full Adder cells in Numerous Logic Styles," in *2022 International Conference on Electronics and Renewable Systems (ICEARS)*, IEEE, Mar. 2022, pp. 90–96, doi: 10.1109/ICEARS53579.2022.9751808.
- [20] V. M. B, S. K. N, and S. D. H, "XNOR-XOR based Full Adder Using Double Pass Transistor Logic," in *2023 International Conference on Advances in Electronics, Communication, Computing and Intelligent Information Systems (ICAECIS)*, IEEE, Apr. 2023, pp. 344–348, doi: 10.1109/ICAECIS58353.2023.10169925.
- [21] V. Dhandapani and S. Ramachandran, "Area and power efficient DCT architecture for image compression," *EURASIP Journal on Advances in Signal Processing*, no. 1, pp. 1–9, Dec. 2014, doi: 10.1186/1687-6180-2014-180.
- [22] H. Naseri and S. Timarchi, "Low-Power and Fast Full Adder by Exploring New XOR and XNOR Gates," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 26, no. 8, pp. 1481–1493, Aug. 2018, doi: 10.1109/TVLSI.2018.2820999.
- [23] K. B. Bhavya and A. M. V. Prakash, "Design and Analysis of Hybrid GDI-PTL Based Wallace Tree Multipliers for Low-Power DCT Architecture in Image Compression," *EAI Endorsed Trans. Internet Things*, vol. 11, Mar. 2026, doi: 10.4108/eetiot.10362.
- [24] S. Tongya, R. Sharma, and G. Sujediya, "Design and analysis of multiplier using GDI and CMOS technology," *International Journal of Innovative Research in Science, Engineering and Technology*, vol. 8, no. 8, 2019, doi: 10.15680/IJIRSET.2019.0808022.
- [25] J. Dhanasekar and M. I. Niranjana, "Design and Analysis of Multipliers using Hybrid Full Adder," in *2023 7th International Conference on Electronics, Communication and Aerospace Technology (ICECA)*, IEEE, Nov. 2023, pp. 383–387, doi: 10.1109/ICECA58529.2023.10395380.
- [26] R. . Nithyashree, S. Afreen, and S. Tantry, "Analysis of various Approximate adders in Ripple Carry Adder design," in *2023 4th IEEE Global Conference for Advancement in Technology (GCAT)*, IEEE, Oct. 2023, pp. 1–5, doi:

- 10.1109/GCAT59970.2023.10353293.
- [27] F. Sabetzadeh, M. H. Moaiyeri, and M. Ahmadinejad, "A Majority-Based Imprecise Multiplier for Ultra-Efficient Approximate Image Multiplication," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 66, no. 11, pp. 4200–4208, Nov. 2019, doi: 10.1109/TCSI.2019.2918241.
- [28] J. Li and Z. Liu, "Efficient compression algorithm using learning networks for remote sensing images," *Applied Soft Computing*, vol. 100, p. 106987, Mar. 2021, doi: 10.1016/j.asoc.2020.106987.
- [29] T. Mendez, V. Kedlaya K, D. Nayak, H. S. Mruthyunjaya, and S. G. Nayak, "A Novel ASIC Implementation of Two-Dimensional Image Compression Using Improved B.G. Lee Algorithm," *Applied Sciences*, vol. 13, no. 16, pp. 1–18, Aug. 2023, doi: 10.3390/app13169094.
- [30] K. K. Senthilkumar, K. Kunaraj, and R. Seshasayanan, "Implementation of computation-reduced DCT using a novel method," *EURASIP Journal on Image and Video Processing*, no. 1, pp. 1–18, Dec. 2015, doi: 10.1186/s13640-015-0088-z.

BIOGRAPHIES OF AUTHORS



Bhavya K B    is an assistant professor at the Department of Electronics and Communication Engineering, Bangalore Institute of Technology, VTU, Bengaluru, where she has been a faculty member since 2017. She graduated with a first-class honours B.Eng. degree in Electronics and Communication Engineering from Visvesvaraya Technological University (VTU), Belagavi, India, in 2012, and an M.Tech. in Signal Processing and VLSI from Jain University, India, in 2014. Currently, she is a research scholar at BIT, VTU, Bengaluru, India. Her research interests are primarily in the area of signal processing, multimedia, low-power VLSI, and biomedical signal processing engineering. She can be contacted at email: bhavyakb@bit-bangalore.edu.in.



Dr. A M Vijaya Prakash    is principal, professor, and head (VLSI Design and Technology) at Bangalore Institute of Technology, VTU, Bengaluru. He holds an M.E., Ph.D., MISTE, and FIE and completed his Ph.D. at Dr. MGR University in the year 2012 in the field of Low-Power VLSI Design. He is a recognized as research guide under Visveswaraya Technological University. Five students are awarded Ph.D. and two number of students are pursuing Ph.D. under his guidance. His major research interests include VLSI design, image processing, and video processing. He has published 48 number of papers in recognized journals/conferences. He can be contacted at email: am_vprakash@bit-bangalore.edu.in.