

A diverse high-performance platform for Advanced Driver Assistance System (ADAS) applications



Prashanth Viswanath
Kedar Chitnis
Pramod Swami
Mihir Mody
Sujith Shivalingappa
Soyeb Nagori
Manu Mathew
Kumar Desappan
Shyam Jagannathan
Deepak Poddar
Anshu Jain
Hrushikesh Garud
Vikram Appia
Mayank Mangla
Shashank Dabral

Texas Instruments



Overview

Advanced driver assistance systems (ADAS) are becoming increasingly popular. ADAS applications such as lane departure warning (LDW), forward collision warning (FCW), automatic cruise control (ACC), auto emergency braking (AEB) and surround view (SV) that were present only in luxury vehicles in the past have trickled down to entry- and mid-level vehicles. Many of these applications are also mandated by safety authorities such as European New Car Assessment Program (Euro NCAP) and National Highway Traffic Safety Administration (NHTSA). In order to make these applications affordable in entry- and mid-level vehicles, it is important to have a cost-effective, yet high-performance and low-power solution. Texas Instruments (TI's) TDA3x is an ideal platform to address these needs. In this paper we will illustrate the mapping of multiple algorithms such as SV, LDW, object detection (OD), structure from motion (SFM) and camera-monitor systems (CMS) to the TDA3x device, thereby demonstrating its computing capabilities. We also share the performance for these embedded vision applications, showing that TDA3x is an excellent high-performance device for ADAS applications.

1. Introduction

In the automotive space, the demand for ADAS technology has increased as mobility has come to be a basic need in today's life. Approximately 1.25 million people died in road accidents around the globe in 2013^[2]. Pedestrians, cyclists and motorcyclists comprise half of the road traffic deaths, and motor vehicle crashes are ranked number nine among top ten leading causes of death in the world^[5]. These statistics push automobile manufacturers to ensure higher safety standards in their vehicles. The European New Car Assessment Program (Euro NCAP) and National Highway Traffic Safety Administration (NHTSA) provide safety ratings to new cars based on the safety systems that are in place. Euro NCAP^[6] provides better star rating for cars equipped with AEB, FCA, LKA and other ADAS applications, which ensures higher safety for on-road vehicles, pedestrians, cyclists and motorcyclists.

ADAS applications can be based upon various sensor systems such as radar, camera, LiDAR and ultrasound^[7]. They can also integrate and use external information sources such as global positioning systems, car data networks and vehicle-to-vehicle or vehicle-to-infrastructure communication systems to efficiently and accurately achieve desired goals. While different sensor modalities have varying performance based on different environmental conditions and applications, camera sensors are emerging as a key differentiator by car manufacturers. Camera-based ADAS systems use various computer vision (CV) technologies to perform real-time driving situation analysis and provide warning to the driver. The advantages of camera-based ADAS include reliability and robustness under difficult real life scenarios, and ability to support multiple varied applications such as traffic sign recognition (TSR), traffic light detection, lane and obstacle detection.

To enable different safety aspects of ADAS, camera-based systems are deployed in front, back and surround view^[8]. The front camera systems are used for applications such as AEB and FCW. The rear view and surround view systems are used for park assist and cross traffic alert applications. Front camera systems can use mono or stereo camera setup. Stereo camera is useful to obtain 3-D information by generating disparity. However, stereo camera systems are more expensive compared to mono camera systems. Structure from Motion (SfM) technology^{[17][11]}, which enables a single moving camera to obtain depth, is being widely researched for its capabilities in ADAS applications. Surround view systems use multiple cameras (four to six) placed around the car. The feed from multiple cameras are re-mapped and stitched to provide a 360-degree view to the driver. Also, analytics are performed on these images to alert the driver. Recently, Camera Mirror Systems (CMS) are increasingly replacing mirrors in mid-/high-end cars. In CMS systems the side and rear

view mirrors are replaced by cameras, and the camera feed is displayed to the driver via display panels (typically OLED display panels). Cameras with wide angle field of view can reduce the number of blind spots for the driver. Sophisticated features like wide dynamic range (WDR)^[15] and noise filter allow the system to be used in variety of lighting conditions including low-light, high-glare scenarios. Due to the low surface area of the camera lens versus a conventional mirror, a CMS system is less susceptible to the effects of dust and rain. CMS systems also have added advantage of reducing wind drag and thus aiding fuel efficiency. Finally, the CMS opens the possibility of running vision analytics on them^[12]. Figure 1 shows the flow for different ADAS applications.

In order to fully utilize the capabilities of camera-based systems for multiple applications, it is imperative to have a high-performance, low-power embedded processor that is capable of analyzing data from multiple cameras in real time. In order to solve this problem, Texas Instruments

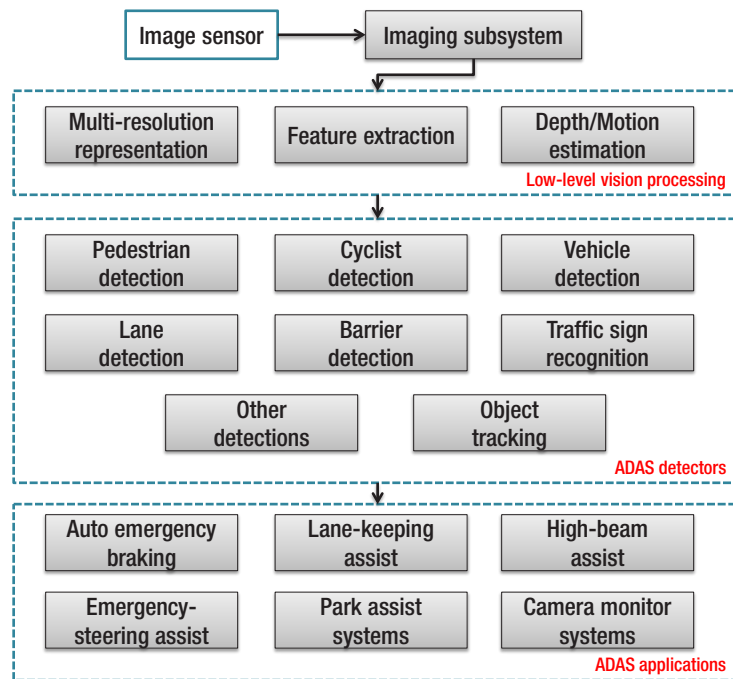


Figure 1: Flow chart of ADAS applications

(TI) has developed a family of System-on-Chip (SoC) processors that integrate heterogeneous compute architectures like general-purpose processor (GPP), digital signal processor (DSP), single-instruction multiple data (SIMD) processor and hardware accelerators (HWA) to satisfy the compute requirements and still meet the area and power specifications. The rest of the paper is organized as follows: Section 2 provides an introduction to a high-performance, low area and power, third generation of SoC solution from TI called Texas Instruments Driver Assist 3x (TDA3x), Section 3 illustrates different applications such as LDW, OD, SFM, SV, CMS and their mapping to the TDA3x platform. Section 4 shows the results of our implementation and the performance data and Section 5 provides the conclusion.

2. TDA3x introduction

The TDA3x SoC^[4] has a heterogeneous and scalable architecture that includes a dual-core ARM[®] Cortexv-M4, dual-core C66x DSP and single-

core Embedded Vision Engine (EVE) for vector processing, as shown in Figure 2. It integrates hardware for camera capture, image signal processor (ISP) and display sub-system resulting in better video quality at lower power. It also contains large on-chip random access memory (RAM), a rich set of input/output peripherals for connectivity, and a safety mechanism for automotive market. There are three types of programmable cores in the TDA3x SoC: GPP, DSP, and EVE.

2.1 General-Purpose Processor (GPP)

The dual-core ARM Cortex-M4 CPU, running at 212.8 MHz, serves as the general-purpose processor in the TDA3x processor^[1]. The M4 cores deliver efficient control and processing camera stream.

2.2 Digital Signal Processor (DSP)

The TDA3x SoC contains a dual-core C66x DSP. The C66x DSP^[3] is a floating-point Very

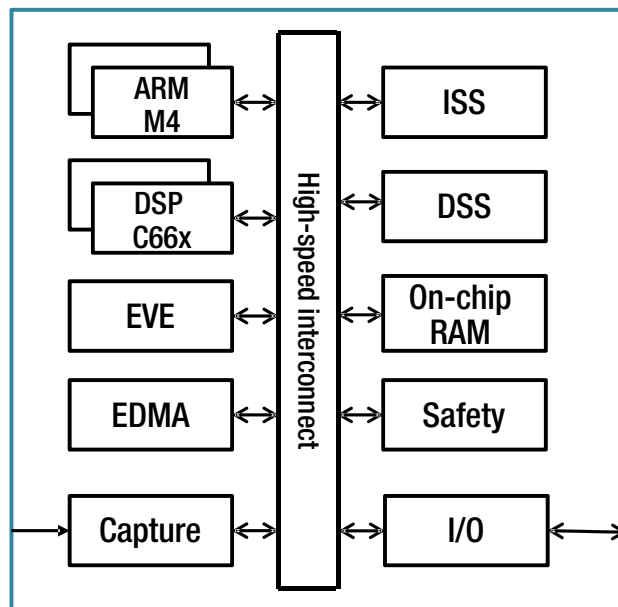


Figure 2: TDA3x SoC block diagram

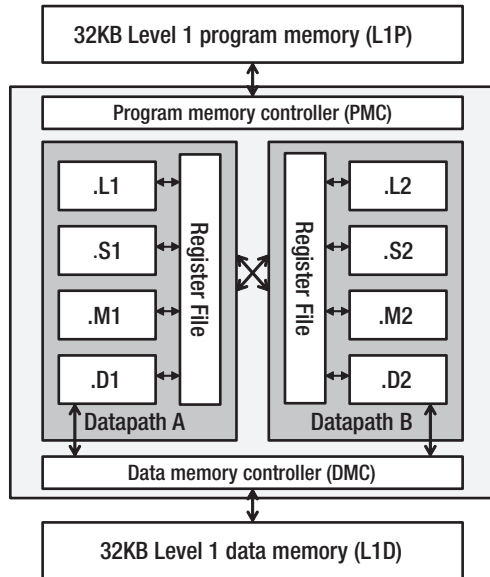


Figure 3: C66x processor block diagram

Long Instruction Word VLIW architecture with 8 functional units (2 multipliers and 6 arithmetic units) that operate in parallel, as shown in Figure 3. It comprises of 64 general-purpose 32-bit registers shared by all eight functional units. There are four arithmetic units .L1/.L2, .S1/.S2, two multiplier units for .M1/.M2 and two data load and store units, .D1/.D2. Each C66x DSP core has configurable 32 KB of L1 data cache, 32 KB of L1 instruction cache and 288 KB of unified L2 data/instruction memory.

2.3. Embedded Vision Engine (EVE)

TI's TDA3x contains a single-core EVE, a fully programmable accelerator specifically to enable the processing, latency and reliability needs found in computer vision applications. The EVE includes one 32-bit Application-Specific RISC Processor (ARP32) and one 512-bit vector coprocessor (VCOP) with built-in mechanisms and unique vision-specialized instructions for concurrent, low overhead processing. The VCOP is a dual 8-way SIMD engine with built-in loop control and address generation.

It has certain special properties such as transpose store, de-interleave load and interleaving store. The VCOP also has specialized pipelines for accelerating table look-up and histograms^[13]. Figure 4 shows the block diagram of EVE processor.

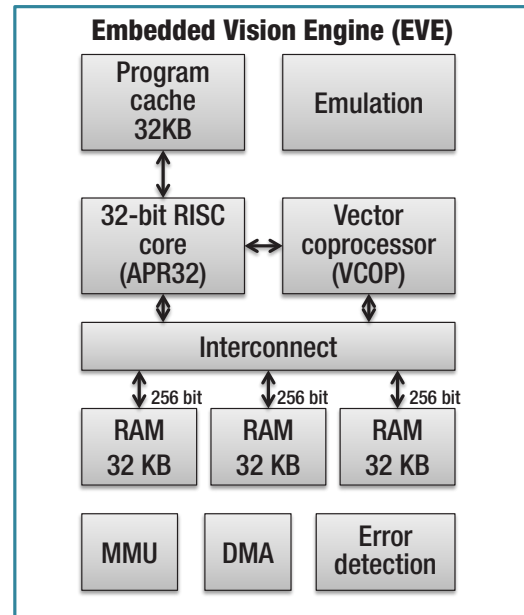


Figure 4: EVE processor block diagram

3. Applications and system partitioning

3.1. System partitioning

A computer vision application can be roughly categorized into three types of processing: Low-level, mid-level and high-level processing. The low-level processing functions include pixel-processing operations, where the main focus is to extract key properties such as edges and corners, and to form robust features. The mid-level processing functions include feature detection, analysis, matching and tracking. High-level processing is the stage where heuristics are applied to make meaningful decisions by using data generated by low- and mid-level

processing. The EVE architecture is an excellent match for low-level and mid-level vision-processing functions due its number-crunching capability. The C66x DSP with program and data caches enables mix of control as well as data-processing capabilities and is well suited for mid- and high-level vision-processing functions. High-level OS (or RTOS) runs on ARM as the main controller and does I/O with real world.

3.2. Object detection and traffic sign recognition

The object detection algorithm consists of low-, mid-, and high-level processing functions, and are mapped across the EVE and DSP cores as shown in Figure 5. As EVE is suitable for low- and mid-level processing, stages such as gradient computation, orientation binning and histogram equalization are mapped to EVE while the classification stage is mapped to the C66x DSP.

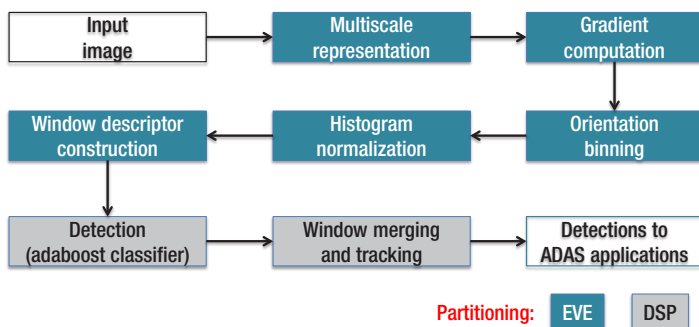


Figure 5: Object detection algorithm partitioning

3.2.1 Gradient computation on EVE

Gradient computation is one of the most commonly used operations in the feature detection stage of various algorithms such as histogram of gradients (HoG)^[9] and ORB^[20]. The gradient is calculated by finding the absolute difference of pixels in horizontal and vertical direction and adding both providing

magnitude of gradient. Figure 6 shows optimized code written in kernel-C (C-like language for EVE) for gradient magnitude computation. Each VCOP computation instruction/line in Figure 6 operates on eight elements. VCOP has two 256 bit functional units each and can operate on eight data elements in parallel. Two instruction/lines can be executed in a cycle. Address computations are performed by dedicated units so it can happen in parallel with core computation. Loop counters are managed by nested loop controller of VCOP, and does not add any overhead. Data load and store instruction can be hidden by compute cycles. The loop in Figure 6 takes just 4 cycles per iteration (generating output for 16 pixel locations in parallel), resulting in 64 times faster performance.

Program: Gradient Magnitude

```

Z = 0;
for (I1 = 0; I1 < height; I1++) {
  for (I2 = 0; I2 < (width/16); I2++) {
    // Separate Address generation hardware
    Addr1 = I1*pitch*ELEMSZ + I2*VECTORSZ*2;
    Addr2 = I1*width*ELEMSZ*2 + I2*VECTORSZ*2*2;
    Addr3 = I1*width*ELEMSZ*2 + I2*VECTORSZ*2*2;
    // Data Load for dual SIMD in VCOP
    (VinT1,VinT2) = (pIn+1)[Addr1].deinterleave();
    (VinL1,VinL2) = (pIn+pitch)[Addr1].deinterleave();
    (VinR1,VinR2) = (pIn+pitch+2)[Addr1].deinterleave();
    (VinB1,VinB2) = (pIn+2*pitch+1)[Addr1].deinterleave();
    // Vector Computation
    VgX_1 = VinR1 - VinL1;
    VgY_1 = VinB1 - VinT1;
    VgX_2 = VinR2 - VinL2;
    VgY_2 = VinB2 - VinT2;
    Vmag1 = abs(VgX_1);
    Vmag2 = abs(VgX_2);
    Vmag1 += abs(VgY_1-Z);
    Vmag2 += abs(VgY_2-Z);
    // Data Store from dual SMID in VCOP
    pGradX[Addr2].interleave() = (VgX_1,VgX_2);
    pGradY[Addr2].interleave() = (VgY_1,VgY_2);
    pMag[Addr3].interleave() = (Vmag1,Vmag2);
  }
}
  
```

Figure 6: Gradient magnitude computation on EVE

3.2.2 Adaboost Classification on C66x DSP

Adaboost classifier uses a set of simple decision trees whose individual classification accuracy is

slightly more than fifty percent^[18]. By combining the response of several such simple decision trees, a strong classifier can be constructed without the need for sophisticated classifier engine as shown in Figure 7. Each individual tree is comprised

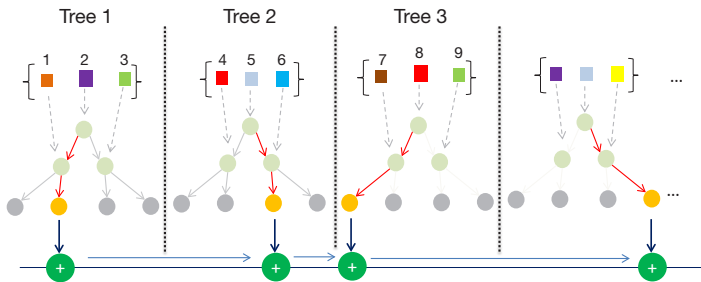


Figure 7: Adaboost classifier diagram

of three nodes and four leaves. Nodes are the locations where an input value is compared against a predetermined threshold. Depending on the comparison result, the tree is traversed left or right until it reaches one of the four possible leaf values. The tree structure, threshold values, leaf values and even the offsets from which input has to be read is predetermined during the training stage. The final leaf value or responses of each tree is accumulated. The accumulated response is compared against a cascade threshold which finally classifies the object. This algorithm is data bound with three thresholds, three offsets, three inputs and four leaf values read for each tree. Assuming that all data is 16 bit, accessing the inputs, off-sets, thresholds, leaf values one at a time will be inefficient on a C66x

DSP with 64-bit data paths. As the C66x DSP supports SIMD processing for fixed-point data, the first step is to load and process four 16-bit data in a single cycle. TI provides intrinsic instructions which can be used to perform SIMD operations. As input data tested at each node is fetched from a sparse offset in the memory, software can perform SIMD loads of the predetermined offsets, thresholds and leaf values stored contiguously in the memory^[16].

3.3. Lane Departure Warning (LDW)

The lane departure warning (LDW) algorithm consists of low- and mid-level processing functions. Due to the presence of a single EVE on the TDA3x, the LDW is mapped to the C66x DSP. Also, since the LDW algorithm uses canny-edge detection that includes the edge-relaxation stage which cannot be made block based, and due to the limited internal memory of EVE, this algorithm is mapped to the C66x DSP for simplicity of design. The block diagram of LDW is shown in Figure 8. The LDW

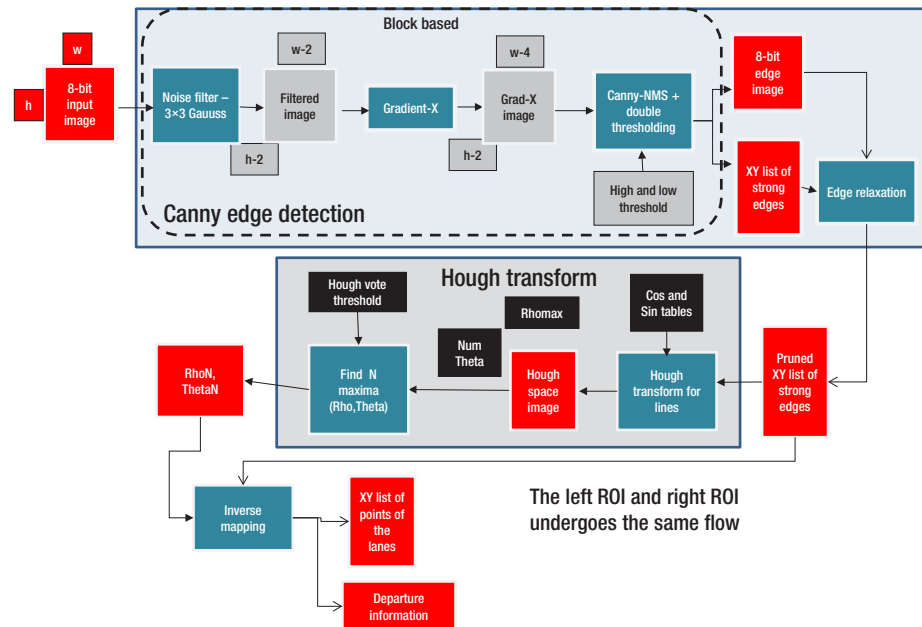


Figure 8: LDW algorithm block diagram

algorithm is purely image based and uses simple processing functions such as canny-edge detection and Hough transform for lines to detect the lanes. Algorithmic enhancements and simplifications such as Intelligent ROI definition, computation of horizontal gradient only, detection of inner/outer edge and curve detection using Hough transform for lines are performed. More details of the algorithm implementation can be found in Figure 17 on page 12. Dataflow optimization techniques such as use of direct memory access (DMA) to transfer smaller blocks of data into L1D, ping-pong data processing are employed to reduce the DDR bandwidth.

3.4. Structure From Motion (SFM)

Structure from motion (SFM) is a key algorithm which enables computation of depth using a single camera which is moving^{[17][11]}. The key components of SFM are sparse optical flow (SOF) and triangulation. Optical flow estimates pixel motion between two temporally ordered images. Lucas Kanade (LK)^[14]-based SOF is widely used for these purposes. Figure 9 shows the various

modules involved in SOF. The SOF algorithm is implemented on the EVE engine of the TDA3x. Although SOF operates on sparse points which is not typically suitable for EVE, the algorithm is designed to operate optimally by operating on multiple sparse points together and also utilizing the DMA engine to organize data suitably, thereby utilizing the SIMD capability of EVE. Special instructions of EVE such as collated-store and scatter help save computation. Collate-store collects all the converged points and then further computation is performed for only those points. Scatter is used to later revert the results back to its original location. Once the SOF algorithm provides reliable optical flow tracks for various key points, triangulation is performed on the C66x DSP to obtain the 3-D point cloud.

3.5. 3D Surround View (SV) system

Surround view (SV) systems are becoming popular in entry- and mid-range cars^[21]. 2D SV systems provide a stationary top-down view of the surroundings from above the vehicle, whereas 3D

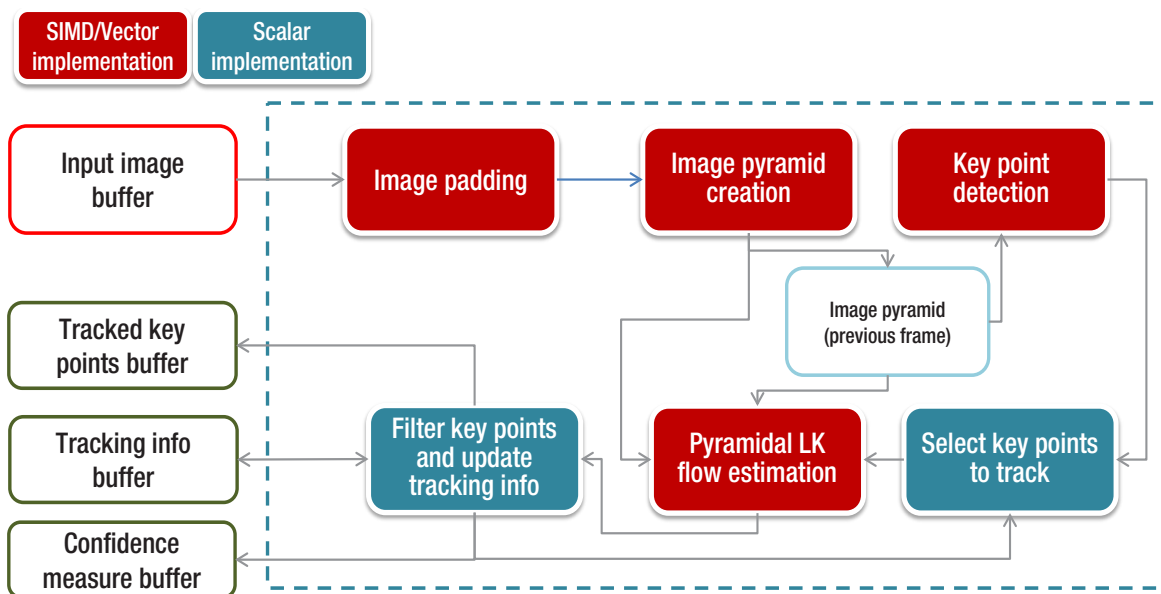


Figure 9: Sparse optical flow block diagram



Figure 10: 2D surround view system output



Figure 11: 3D surround view system output

SV systems provide rendering capabilities of the surroundings of vehicle from any virtual viewpoint and transitions between viewpoints. Examples of 2D SV and 3D SV are shown in Figures 10 and 11. An SV system is constructed by stitching multiple video cameras placed around the periphery of the vehicle as shown in Figure 12. Typically, a

dedicated graphics processing unit (GPU) processor would be employed for composition of 3D SV. Since the TDA3x does not have a GPU, the 3D SV is implemented using the combination of a lens distortion correction (LDC) accelerator and a C66x DSP. The GPU typically stores the entire representation of the 3D world, allowing the user to




Figure 12: Surround view algorithm flow

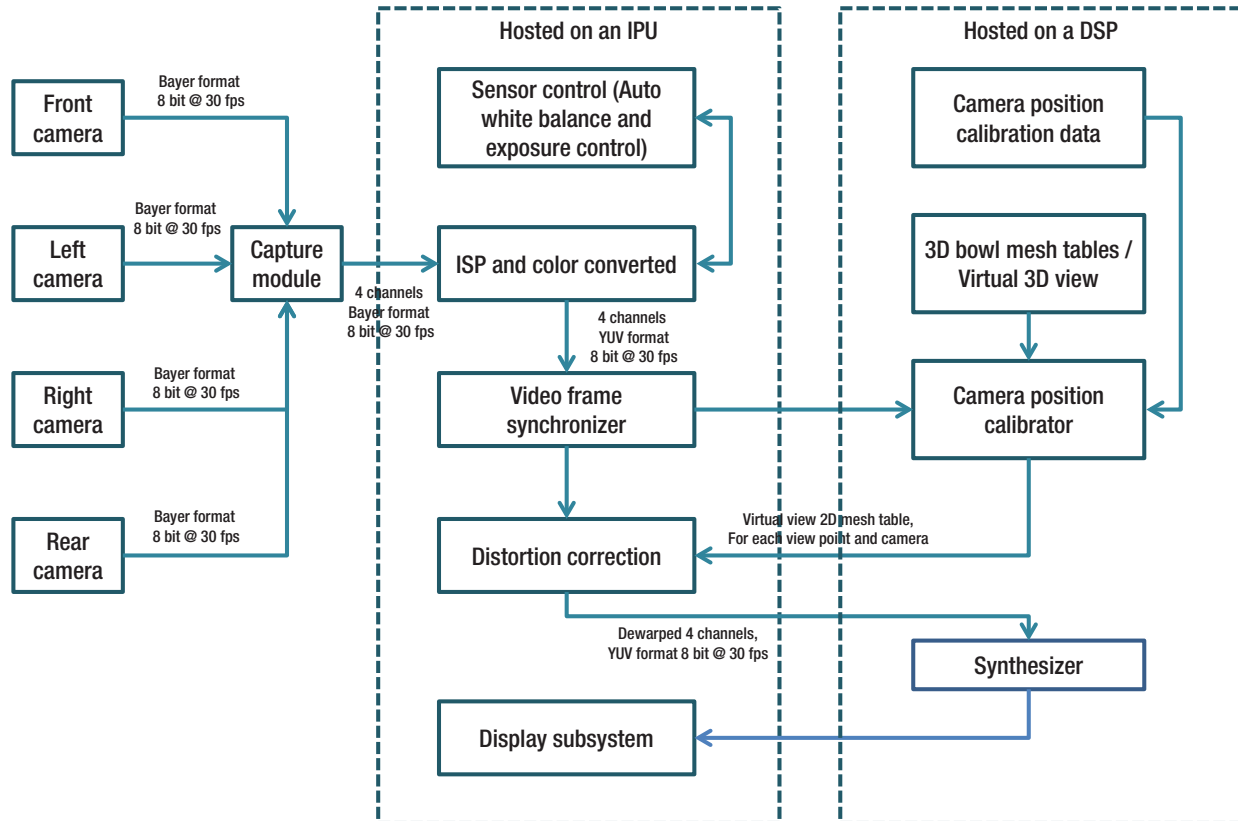


Figure 13: Surround view data flow

change viewpoints. However, this can be optimized by projecting only those 3D points that are in the visible region. Distortion correction is required to correct the fish-eye distortion present in the image sensors used in SV systems. The ISP of the TDA3x SoC has a robust LDC accelerator which performs the distortion correction. In order to create multiple viewpoints, a minimized 3D world map for all the cameras and viewpoints are generated and these viewpoints are stored in non-volatile memory. This can be done both offline and once during the camera setup. When the system boots up, for all the valid viewpoints, 3D world map is read and associated LDC mesh table is generated. Then, these outputs from the LDC are stitched together to obtain the 360°-view for any viewpoint. The data flow for the same is shown in Figure 13.

3.6. Camera Monitor System (CMS)

Figure 14 shows the placement of the cameras in a typical CMS. The algorithm processing involved in a CMS system and its partitioning in TDA3x SoC is shown in Figure 15 on the following page. CMOS sensors are used to capture the scene. A frame-rate of 60 fps is typically utilized to reduce




Figure 14: CMS camera placement illustration

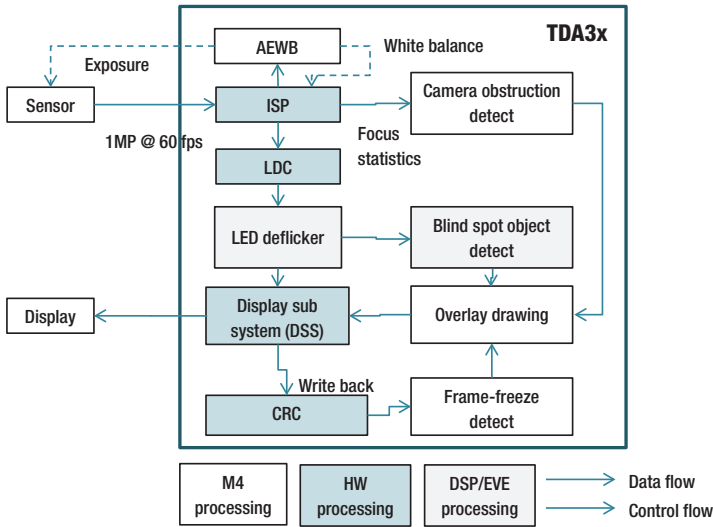


Figure 15: Algorithm and data flow in CMS

latency of the scene as viewed by the car driver. The data format of CMOS sensors is typically Bayer raw data and it is passed through many stages of the ISP before it is converted to viewable video data as shown in Figure 16. Key features in the ISP are: spatial noise filter, which helps to improve the image quality in low-light conditions, and wide dynamic range, which increases the dynamic range of the scene so that bright areas are not clipped and at the same time details in shadow (or darker) regions are visible. This allows the CMS system to operate in a variety of lighting conditions. ISP in the TDA3x also outputs auto white balance (AWB) [10] and auto exposure (AE) statistics which are used by the AEWB algorithm to dynamically adjust the sensor exposure and scene white balance to adapt the cameras settings, dynamically changing

lighting conditions. Additionally, focus statistics are output by the ISP indicating the degree to which the camera is in focus. When the camera lens is obstructed by dirt or water, the scene will not be in focus. Due to the safety-critical nature of the CMS application, it is important to detect such scenarios. Focus statistics can be used in algorithms to detect such events, thereby warning the user of suboptimal scene capture. The hardware LDC module is used to adjust for lens distortion correction due to wide-angle field of view.

A common problem associated with cameras for visual systems like CMS is LED flickering. LEDs are commonly used in car headlights, traffic signals and traffic signs. LEDs are typically pulsed light sources. However, due to persistence of vision, our eyes cannot see the LED flicker. However, camera sensors, especially when they operate at low exposure time due to bright lighting conditions, could capture LED pulses in one frame and miss the LED pulse in the next frame causing an unpleasing and unnatural flicker-like effect. Worst case, it could happen that the LED pulse is not captured at all, say a red light or car headlight, thus giving a dangerous false scene representation to the user. A deflicker algorithm is typically used to eliminate the flicker due to LED lights. This is a pixel-processing-intensive algorithm and is typically run on DSP/EVE. After the LED deflicker algorithm, the scene is displayed on the display through the display sub-system (DSS). A key safety measure in a CMS system is informing




Figure 16: ISP data flow in CMS

the user in case of a frame freeze scenario. Since the user is not constantly looking at the mirror, it could happen that due to hardware (HW) or software (SW) failure, the data that is displayed on the screen is frozen with the same frame repeating. This can cause a hazardous situation to the road users. In the TDA3x, this can be detected by using the DSS to write back the pixel data that is being displayed and then computing a CRC signature for the frame using a HW CRC module. If the CRC signature matches for a sequence of consecutive frames, then it implies that there is frame freeze somewhere in the system and a warning is given to the user or the display is blanked out. Additional analytics algorithms like object detect can be run in the blind spot to notify the driver.

4. Results and analysis

In this section, we provide details of system partitioning and performance of multiple applications executing on the TDA3x SoC. TI's TDA3x EVM is used as the experimental platform. In order to showcase our algorithms, we captured multiple scenes with various camera sensors placed around the car. The video sequence contained urban roads with pedestrians, vehicles, traffic signs, lane marking, traffic lights and parking lot scenarios. This video sequence is then decoded via an HDMI player and fed to the TDA3x EVM as shown in Figure 17. The algorithms then utilize all of the available

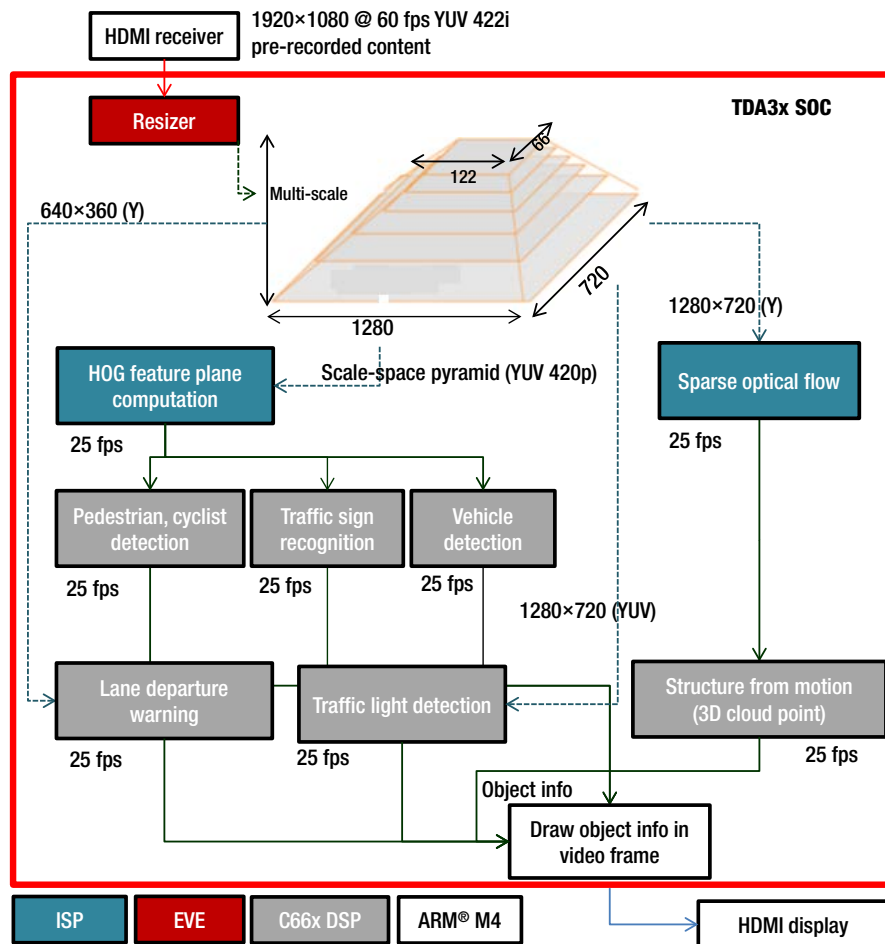


Figure 17: Algorithm partitioning for front camera applications on TDA3x SoC

compute blocks such as ISP, EVE, DSP and ARM Cortex-M4 to perform various functions such as OD, LDW, TSR, SFM, SV and CMS. The output of these algorithms is provided to the ARM Cortex-M4 to draw these markings in the original video and send out the annotated video to the HDMI display. An LCD display is used to watch the video along with object markings to confirm the expected behavior of algorithms. The configuration parameters of these algorithms are listed in Table 1.

| Algorithm | Frame rate | Configuration details on TDA3x SoC |
|-----------------------------------|------------|--|
| Vehicle, pedestrian, cycle detect | 25 | Resolution = 1280 × 720, multi-scale Minimum object size = 32 × 64 |
| Traffic sign recognition | 25 | Resolution = 1280 × 720, multi-scale Minimum object size = 32 × 32 |
| Traffic light recognition | 25 | Resolution = 1280 × 720 Radii range = 8 |
| Lane depart warning | 25 | Resolution = 640 × 360 Number of lanes detected = 2 |
| Structure from motion | 25 | Resolution = 1280 × 720 Number of SOF tracks = 1K points 3D cloud points generated = 800 |
| Surround view system | 30 | Input resolution = 4 channels of 1280 × 800 Output resolution = 752 × 1008 |
| Camera mirror system | 60 | Number of video channels = 1 Input resolution = 1280 × 800 |

Table 1: Algorithm configurations

In the case of front camera applications, the capture is configured for 25 frames per second (fps). As a first step, multiple scales of the input frame are created by using resizer functionality in the ISP. The scale space pyramid is useful to detect objects of different sizes with a fixed-size template. For every relevant pixel in these scales, histogram of oriented gradients (HOG)^[9] signatures are formed. This module involves intensive calculations at the pixel level and hence executed on the EVE. After formation of the HOG feature plane, the EVE

runs the SOF algorithm and the C66x DSP1 runs the adaboost classifier stage of object detection algorithm. The classifier is executed separately for each object category such as pedestrians, vehicles, cyclists and traffic signs. From the scale space pyramid, 640 × 360 and 1280 × 720 scale is fed to DSP2 on which lane detection and traffic light recognition algorithms are run. After completion of SOF, the EVE sends the optical flow tracks to DSP2 to perform triangulation to obtain 3D location of key points in the frame, thereby helping to identify distance of various objects in the scene. In this setup, the ARM Cortex-M4 manages the capture and display device, feeds the data to the ISP and collects information from the DSPs before annotating the objects and displaying them.

For SV application, four channels of video of resolution 1280 × 800 at 30 fps are captured from RAW video sensors in Bayer format, which is supported by the ISP. The ISP then converts the Bayer format data to YUV format for further processing. Auto white balance and exposure-control algorithms ensure each video source is photometrically aligned. Then, the camera calibrator will generate the required mesh table for distortion correction based on the view point and distortion correction is performed. Synthesizer will then receive the corrected images and stitch to form the SV output with a resolution of 752 × 1008 at 30 fps.

In the case of CMS, each camera input operates on one TDA3x SoC. Each channel of video of resolution 1280 × 800 at 60 fps are captured from RAW video sensors (Bayer format) which is supported by the ISP. The ISP then converts the Bayer format data to YUV format, as shown in Figure 16. Algorithms such as OD are run for blind spot detection. Also, a deflicker algorithm is run to remove any LED flicker-related issues, before displaying it to the driver.

| Algorithms | DSP1 utilization | DSP2 utilization | EVE utilization | ARM Cortex-M4 utilization | Frame rate (fps) |
|------------------------|------------------|------------------|-----------------|---------------------------|------------------|
| Front camera analytics | 53% | 66% | 79% | 33% | 25 |
| Surround view system | 45% | 0% | 0% | 44% | 30 |
| Camera mirror system | 68% | 20% | 40% | 30% | 60 |

Table 2: Performance estimates for different applications on the TDA3x SoC

Table 2 shows the loading of the various processors of the TDA3x SoC while running these algorithms.

For front-camera applications, 53 percent of DSP1, 66 percent of DSP2, 79 percent of EVE and 33 percent of one ARM Cortex-M4 are utilized.

For SV application, 45 percent of DSP1 is utilized and 44 percent of one ARM Cortex-M4 is utilized.

The unused C66x DSP and EVE can be used to run analytics on the SV output if needed. For CMS application, 68 percent of DSP1 is consumed to run the deflicker algorithm and 20 percent of DSP2 and 40 percent of EVE is used for the blind spot detection algorithm.

5. Conclusion

ADAS applications require high-performance, low-power and low-area solutions. In this paper, we have presented one such solution based on Texas Instruments' TDA3x device. The paper provided insight into key algorithms of ADAS such as front camera, surround view and camera monitor systems and presented the system partitioning of these algorithms across multiple cores present in the TDA3x SoC and their performance has been provided. Additionally, it has shown that TI's TDA3x platform is able to generate 3D SV efficiently, without a GPU. We have also shown that the TDA3x platform is able to map various ADAS algorithms and still have headroom for customer's differentiation. For more information on the TDA3x platform visit www.ti.com/TDA.

References

- [1] Cortex-M4: Technical reference manual, ARM Inc., 2010.
- [2] Global Health Observatory (GHO) data, 2013.
- [3] TMS320C66x DSP: User guide, SPRUGW0C, Texas Instruments Inc., 2013.
- [4] TDA3x SoC processors for advanced driver assist systems (ADAS) technical brief, Texas Instruments Inc., 2014.
- [5] The top 10 causes of death, 2014.
- [6] 2020 Roadmap, Rev. 1, European New Car Assessment Programme, 2015.
- [7] Advanced driver assistance system (ADAS) guide 2015, Texas Instruments Inc., 2015. Supplied as material SLYY044A.
- [8] S. Dabral, S. Kamath, V. Appia, M. Mody, B. Zhang, and U. Batur. Trends in camera-based automotive driver assistance systems (ADAS). IEEE 57th International Midwest Symposium on Circuits and Systems (MWSCAS), pages 1110–1115, 2014.
- [9] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), 1:886–893, 2005.
- [10] H. Garud, U. K. Pudipeddi, K. Desappan, and S. Nagori. A fast color constancy scheme for automobile video cameras. International

Conference on Signal Processing and Communications (SPCOM), pages 1–6, 2014.

- [11] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [12] B. Kisaanin and M. Gelautz. Advances in Embedded Computer Vision. *Advances in Computer Vision and Pattern Recognition*. Springer, 2014.
- [13] Z. Lin, J. Sankaran, and T. Flanagan. Empowering automotive with TI's vision accelerationpac, 2013. <http://www.ti.com/lit/wp/spry251/spry251.pdf>.
- [14] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. *International Joint Conference on Artificial Intelligence*, 81:674–679, 1981.
- [15] M. Mody, N. Nandan, H. Sanghvi, R. Allu, and R. Sagar. Flexible wide dynamic range (WDR) processing support in image signal processor (ISP). *IEEE International Conference on Consumer Electronics (ICCE)*, pages 467–479, 2015.
- [16] M. Mody, P. Swami, K. Chitnis, S. Jagannathan, K. De-sappan, A. Jain, D. Poddar, Z. Nikolic, P. Viswanath, M. Mathew, S. Nagori, and H. Garud. High-performance front camera ADAS applications on TI's TDA3x platform. *IEEE 22nd International Conference on High-Performance Computing (HiPC)*, pages 456–463, 2015.
- [17] P. Sturm and B. Triggs. A factorization-based algorithm for multi-image projective structure and motion. *European Conference on Computer Vision (ECCV)*, 2:709–720, 1996.
- [18] P. Viola and M. Jones. Fast and robust classification using asymmetric adaboost and a detector cascade. *Advances in Neural Information Processing System*, 2001.
- [19] P. Viswanath and P. Swami. A robust and real-time image-based lane departure warning system. In *Proceedings*. *IEEE Conference on Consumer Electronics (ICCE)*, 2016.
- [20] P. Viswanath, P. Swami, K. Desappan, A. Jain and A. Pathayapurakkal. Orb in 5 ms: An efficient SIMD friendly implementation. *Computer Vision-ACCV 2014 Workshops*, pages 675–686, 2014.
- [21] B. Zhang, V. Appia, I. Pekkucuksen, A. U. Batur, P. Shastry, Liu, S. Sivasankaran, K. Chitnis, and Y. Liu. A surround video camera solution for embedded systems. *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 676–681, 2014.

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

The platform bar is a trademark of Texas Instruments. All other trademarks are the property of their respective owners.

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, enhancements, improvements and other changes to its semiconductor products and services per JESD46, latest issue, and to discontinue any product or service per JESD48, latest issue. Buyers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All semiconductor products (also referred to herein as "components") are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

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

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

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

Reproduction of significant portions of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

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

Buyer acknowledges and agrees that it is solely responsible for compliance with all legal, regulatory and safety-related requirements concerning its products, and any use of TI components in its applications, notwithstanding any applications-related information or support that may be provided by TI. Buyer represents and agrees that it has all the necessary expertise to create and implement safeguards which anticipate dangerous consequences of failures, monitor failures and their consequences, lessen the likelihood of failures that might cause harm and take appropriate remedial actions. Buyer will fully indemnify TI and its representatives against any damages arising out of the use of any TI components in safety-critical applications.

In some cases, TI components may be promoted specifically to facilitate safety-related applications. With such components, TI's goal is to help enable customers to design and create their own end-product solutions that meet applicable functional safety standards and requirements. Nonetheless, such components are subject to these terms.

No TI components are authorized for use in FDA Class III (or similar life-critical medical equipment) unless authorized officers of the parties have executed a special agreement specifically governing such use.

Only those TI components which TI has specifically designated as military grade or "enhanced plastic" are designed and intended for use in military/aerospace applications or environments. Buyer acknowledges and agrees that any military or aerospace use of TI components which have **not** been so designated is solely at the Buyer's risk, and that Buyer is solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI has specifically designated certain components as meeting ISO/TS16949 requirements, mainly for automotive use. In any case of use of non-designated products, TI will not be responsible for any failure to meet ISO/TS16949.

Products

| | |
|------------------------------|--|
| Audio | www.ti.com/audio |
| Amplifiers | amplifier.ti.com |
| Data Converters | dataconverter.ti.com |
| DLP® Products | www.dlp.com |
| DSP | dsp.ti.com |
| Clocks and Timers | www.ti.com/clocks |
| Interface | interface.ti.com |
| Logic | logic.ti.com |
| Power Mgmt | power.ti.com |
| Microcontrollers | microcontroller.ti.com |
| RFID | www.ti-rfid.com |
| OMAP Applications Processors | www.ti.com/omap |
| Wireless Connectivity | www.ti.com/wirelessconnectivity |

Applications

| | |
|-------------------------------|--|
| Automotive and Transportation | www.ti.com/automotive |
| Communications and Telecom | www.ti.com/communications |
| Computers and Peripherals | www.ti.com/computers |
| Consumer Electronics | www.ti.com/consumer-apps |
| Energy and Lighting | www.ti.com/energy |
| Industrial | www.ti.com/industrial |
| Medical | www.ti.com/medical |
| Security | www.ti.com/security |
| Space, Avionics and Defense | www.ti.com/space-avionics-defense |
| Video and Imaging | www.ti.com/video |

TI E2E Community

e2e.ti.com