

Challenges and Solutions for Emerging Automotive Advanced Driver Assistance Systems

The Trend

A growing trend for improving safety in the automotive industry, is the integration of multiple periphery camera sensors and intelligent image processing technology to enable advanced driver assistance systems (ADAS) such as lane departure warning, collision avoidance, blind spot monitoring, enhanced back-up camera and surround view with object detection. This trend is driven by a common goal of manufacturers to create a completely safe driving environment. For example, Volvo has the “Zero-by-2020” goal states that no one will be killed or injured in a new Volvo car by 2020.



Figure 1: Advanced Driver Assistance Systems

The Challenges

The challenges with these systems lies in identifying a platform that can sustain the processing performance that these compute-intensive applications demand, consume low power to avoid heat dissipation issues, provide a cost effective solution that integrators will opt for, and be offered in a small system footprint so the ‘brains’ of the system may be co-located with the image sensor to fit in the confines of the smart camera (if required). Co-locating image processing at the sensor establishes a system solution in which the same intelligent camera platform can be used at different sites in the vehicle such as in the rear bumper to enable enhanced backup camera with object/person detection, side mirrors to handle blind spot monitoring, behind the rearview mirror to handle forward collision and lane departure warning, and in other periphery cameras used for surround view. Furthermore, this distributed intelligent camera model will not burden the vehicle’s central console unit with additional processing requirements.

Market reports indicate that ADAS will proliferate not only in high-end vehicles, but also in more common lower-class vehicles over the next few years. In addition to the demand for more compute performance per application, there is also a drive to bundle more and more ADAS applications onto the same hardware platform. This begs the question – Will present-day DSP and FPGA solutions measure up?

One of the biggest bottlenecks in DSP algorithm execution is the load placed on external memory to keep up with read and write accesses. Traditional DSPs offer limited parallelism and often require operating at higher and higher clock frequencies in order to meet processing requirements. As clock speeds ramp higher and higher, more power is consumed which generates more heat that must be dissipated. While FPGAs offer a higher degree of parallelism over DSPs, they are generally more difficult to program and often require a RISC for post-processing data. FPGAs have high power consumption, large system footprint and are a more costly solution overall.

The Future of Image Cognition: Multi-Core Parallel Processing

Multi-core parallel processing performed by CogniVue™ Image Cognition Processors (ICPs) excel in all areas where DSPs and FPGAs fall short. ICPs deliver more performance, per area, per milliwatt and are offered at lower price point.

CogniVue APEX™ architecture combines an industry-standard RISC core that manages algorithm execution with a massively parallel single instruction multiple data (SIMD) array processing engine (APU) to execute low level compute intensive parallel operations inherent in image processing and analysis algorithms. In addition to the RISC and APU are stream DMAs devised to ensure efficient data movement, and a sequencer conceived to automatically and efficiently order operations ensuring maximum efficiency. A second RISC core operates independently outside APEX to handle system level program routines.

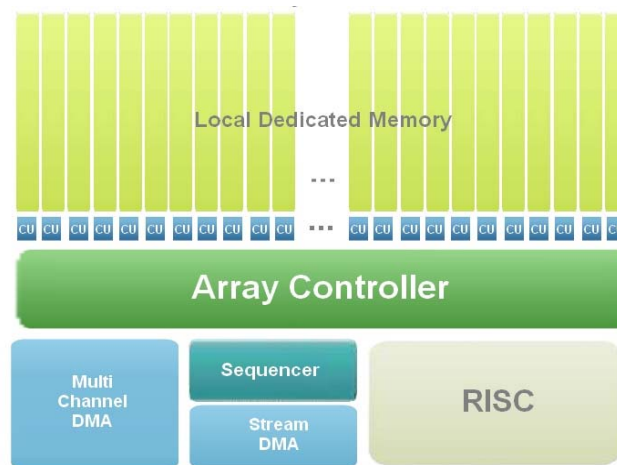


Figure 2: Programmable Parallel Architecture

Within the APU, local dedicated memory is present for each compute element. Image data is fetched from external memory once, streamed into the APU memory where all processing is performed locally before the data is streamed out and stored back in system memory. By co-locating memory with APU elements, the number of external memory accesses is greatly reduced and compute performance is sustained without having to increase clock speed. Furthermore, the APEX processing core is decoupled from the rest of the ICP. This means the APEX core operating frequency is independent of the rest of the SoC and therefore allows the rest of the part to run at lower clock frequencies for more power savings. A smaller system footprint is achieved, which translates to savings in board real estate, by stacking external memory inside the package.

By leveraging a parallel processor core and a unique stream programming-based software paradigm, CogniVue ICPs are capable of scheduling complex vector operations and execute with minimal data movement. These processors enable automatic pipelining of algorithm primitives wherever possible and abstract this complexity through a comprehensive API, thus hiding complexities of system load-balancing and eliminating multi-core synchronization issues from the developer.

A high degree of parallelism with high ALU bandwidth architecture ensures a viable platform providing enough headroom and processing power to execute multiple applications concurrently on the same hardware. A flexible development platform and SDK empowers customers to program APEX and create competitive differentiated applications. This multi-core parallel platform becomes a developer's choice as it not only comes up ahead in performance, power and size, but also assures code reuse in future ICPs thereby guaranteeing reduced development effort and time.