


SOLUTION BRIEF

AUTOMATED DRIVING DOMAIN CONTROLLER WITH VERSAL™ AI EDGE SERIES GEN 2

- Offering a comprehensive processing platform contained within a single adaptive SoC device
- Providing performance leadership with extensive I/O capability, a virtualized software stack, and robust ecosystem
- Includes next-generation AI Engine technology offering up to 3X performance per watt improvement¹

AMD
together we advance_



OVERVIEW

The heterogeneous architecture of the AMD Versal AI Edge Series Gen 2 adaptive SoCs allows for a single-chip solution to tackle all phases of an automated driving system—sense, perceive, plan, and act.

Sensor inputs (e.g., vision, radar/LiDAR, etc.) are ingested via programmable I/O blocks and fed directly to the programmable logic (PL) for low-latency sensor-specific processing. Versal AI Edge Series Gen 2 devices include hardened accelerators for functions like ISP and an ASIL-capable GPU for warping and graphics rendering. The PL offers the most flexible means of implementing innovative sensor fusion algorithms prior to perception/inference processing in the AI Engines.

Behavioral planning is implemented in the application processors targeted for ASIL B (split mode)/ASIL D (lock-step) to assess the environmental perception output and decide what future maneuvers need to be executed. Decisions are redundantly checked, and maneuver actuation is communicated by the ASIL D capable real-time processing domain. The versatility and scalability of the Versal AI Edge Series Gen 2 devices support various L2/L2+ system requirements as well as L3 and L4 systems where redundancy is critical.

HIGHLIGHTS

PROGRAMMABLE I/O FOR SENSOR INGEST AND COMMUNICATIONS

- Customizable interfaces to support evolution in sensor communications
- Off-the-shelf and customizable IP cores for optimized performance and partitioning between software and hardware acceleration
- Scalable sensor interfacing/processing enabling sensor fusion innovation
- Unrivalled system future-readiness with “Hardware OTA” updates

NEXT-GENERATION AI ENGINES FOR NEURAL NETWORK PROCESSING

- New native FP8, FP16, MX6, and MX9 data type support
- Industry-standard framework support
- Local, flexible memory tiles for high-efficiency compute
- Both neural network inference and signal processing workload support

PROGRAMMABLE LOGIC FOR FLEXIBLE, LOW-LATENCY COMPUTE ACCELERATION

- Adaptable on-chip memory hierarchy to match the right memory to the right task
- Ability to adapt to evolving workloads without replacing hardware
- Dynamic hardware reconfiguration support to respond to real-time conditions
- Physical isolation and redundancy for fault resilience and functional safety

HETEROGENEOUS SCALAR COMPUTE IN A UNIFIED DEVELOPMENT TOOL

- CPU cluster structures offer unmatched flexibility in ISO26262 ASIL B/D performance trade-off
- Programmable logic accelerators significantly reduce overall scalar CPU performance needs
- Up to 10X scalar compute vs. previous generation for complex workloads²
- Additional high-speed interfaces including USB 3.2, DisplayPort™ 1.4, multiple PCIe® Gen5x4, and 10G Ethernet

KEY BENEFITS

ADAPTABILITY

Flexibility to meet various vehicle platform ADAS/Automated Driving system and sensor configurations

MODULAR SCALABILITY

Cost-effectively accommodate varying feature sets with a set of adaptable SoCs and IP portability across devices

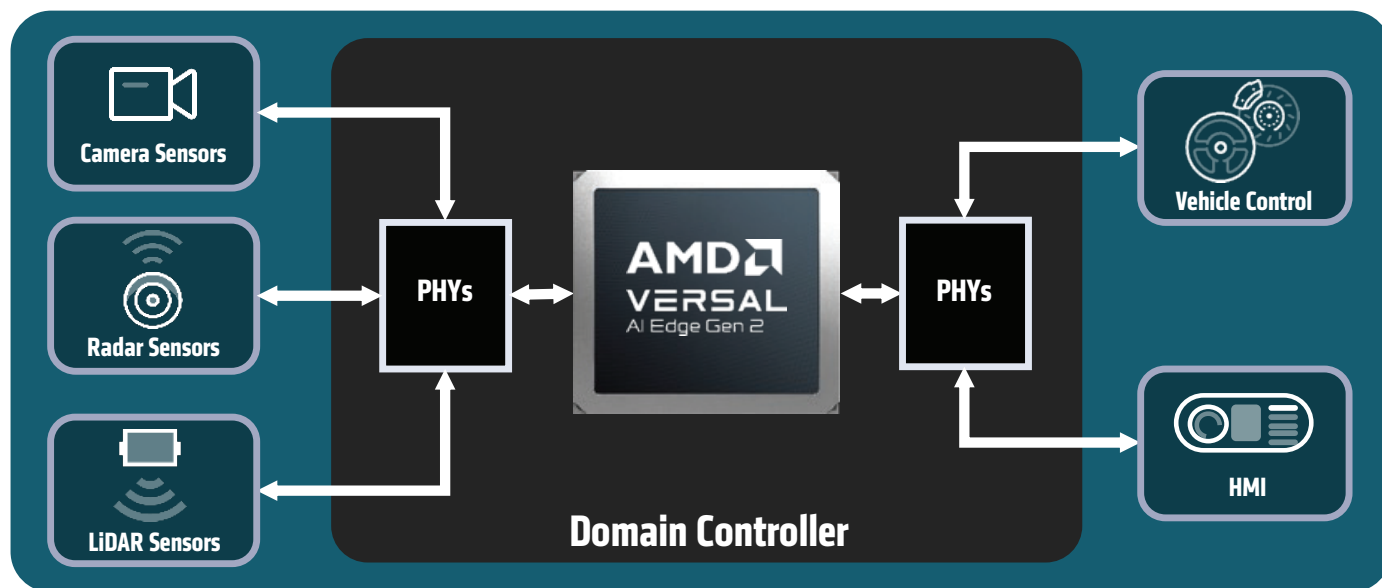
STATE OF THE ART FUNCTIONAL SAFETY

Isolation and decomposition of safety critical elements

OPTIMAL FUNCTIONAL PARTITIONING

High-efficiency and low-latency performance enabled by independent & simultaneous hardware accelerator pipelines

BLOCK DIAGRAM



FEATURES

PLATFORM HIGHLIGHTS	
Flexible Processing Subsystem	<ul style="list-style-type: none"> Octal core Arm® Cortex®-A78AE processors with partitioning capability to satisfy safety and security requirements Deca core Arm Cortex-R52 processors for moderate compute and real-time processing, and system management functionality
Multi-Media Capability	<ul style="list-style-type: none"> Quad core Arm Mali™-G78AE GPU with two slices to enable multi-display use, safety-critical redundancy, and/or graphics overlay demands Encoder & decoder supports HEVC & AVC up to 4K60, 444, 12-bit³
Memory Partitioning and Network on Chip	<ul style="list-style-type: none"> Up to five DDR5/LPDDR5/LPDDR5X memory controllers supporting inline encryption and ECC for up to 170 GB/s bandwidth of memory vs. previous generation,⁴ with inline encryption to secure data Partitionable memory controllers and network on chip to ensure separation
AI Compute	<ul style="list-style-type: none"> Up to 185 INT8 dense TOPS, 370 sparse TOPS⁵ Up to 3X TOPS/watt vs. previous generation¹
Hardened Image Signal Processing	<ul style="list-style-type: none"> Support up to 5 camera streams with real-time streaming or multicasting memory-mapped configuration and throughput of 1.2 Gpix/s per ISP tile Accelerate standard image processing with support for 12+ camera streams and up to 3+ Gpix/s combined throughput

NEXT STEPS

As the pace of innovation continues to accelerate in the automotive industry, the need for high-performance compute, compute acceleration and graphics technologies is increasing. AMD is a leader at this inflection point, with a broad line of high-performance CPUs, GPUs, FPGAs, and adaptive SoCs. For more information, visit the AMD Automotive website: www.amd.com/automotive

To learn more about AMD Versal AI Edge Series Gen 2, visit www.amd.com/versal-ai-edge-gen2

ENDNOTES

1. Based on AMD internal performance and power projections for the AIE-ML v2 compute tile architecture in the Versal AI Edge Series Gen 2 using the MX6 data type, compared to performance specifications and AMD Power Design Manager power results for the AIE-ML compute tile architecture featured in the first-generation Versal AI Edge Series using INT8 data type. Assumptions: 2 row, 8 column sub-arrays. Operating conditions: 1 GHz F_{MAX} , 0.7V AIE operating voltage, 100°C junction temperature, typical process, 60% vector load, % activations = 0 < 10%. Actual performance will vary when final products are released in market. Performance projections as of March 2024. (VER-023)
2. Based on AMD internal pre-silicon performance estimates for combined total DMIPs of the Versal AI Edge Series Gen 2 and Versal Prime Series Gen 2 processing system when configured with 8 Arm Cortex-A78AE applications cores @2.2 GHz and 10 Arm Cortex-R52 real-time cores @1.05 GHz, compared to the published combined total DMIPs of the processing system in the first-generation Versal AI Edge Series and Versal Prime Series. Versal AI Edge Series Gen 2 and Versal Prime Series Gen 2 operating conditions: Highest available speed grade, 0.88V PS operating voltage, split-mode operation, maximum supported operating frequency. First-generation Versal AI Edge Series and Versal Prime Series operating conditions: Highest available speed grade, 0.88V PS operating voltage, maximum supported operating frequency. Actual DMIPs performance will vary when final products are released in market. (VER-027)
3. Video codec acceleration (including at least the HEVC (H.265), H.264, VP9, and AV1 codecs) is subject to and not operable without inclusion/installation of compatible media players. (GD-176)
4. Based on AMD engineering pre-silicon performance estimates for the Versal AI Edge Series Gen 2 2VE3858 device with 5x 32B memory controllers and expected maximum LPDDR5X memory data rate of 8.533 GB/s, compared to an in-production first-generation Versal AI Edge Series VE2802 device with 3x 64b memory controllers operating at the published maximum LPDDR4X memory bandwidth of 102.4 GB/s. Actual memory bandwidth calculations for the Versal AI Edge Series Gen 2 devices are subject to change when final products are released in market. (VER-031)
5. Based on AI Engine clock frequency of 1.25 GHz

DISCLAIMERS

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

COPYRIGHT NOTICE

© 2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Versal, and other designated brands included herein are trademarks of Advanced Micro Devices, Inc. Arm, Cortex, and Mali are trademarks of Arm in the EU and other countries. DisplayPort and the DisplayPort logo are trademarks owned by the Video Electronics Standards Association (VESA®) in the United States and other countries. PCIe and PCI Express are trademarks of PCI-SIG and used under license. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. PID# 242620247-A