

# ALVEO™ MA35D ACCELERATOR

## Unleashing Interactive Streaming at Scale

### OVERVIEW

With the global video market now being dominated by live streaming, a new class of low latency applications are emerging and influencing the cost structure and associated deployment strategy of infrastructure and video processing technologies.

The Alveo™ MA35D media accelerator is ushering in a new era of interactive media applications that scales for cost-, area-, and power-per video channel. Supporting high channel density at up to 32x 1080p60 streams, the accelerator is optimized to reduce rack space, network bandwidth, and power consumption—all while maintaining video quality at ultra-low latency. By helping to reduce CAPEX and OPEX, infrastructure and content providers can cost-effectively scale while ensuring optimal quality of experience.

With a Half-Height, Half-Length (HHHL) PCIe® form factor, the Alveo MA35D card delivers comprehensive support for video codec technologies including H.264, H.265, and the emerging AV1 standard.

### HIGHLIGHTS

#### Video Processing Unit for High Density, Ultra-Low Latency Streaming

- Specialized ASIC architecture reduces cost-, area-, and power-per stream
- Support for AV1, H.264, H.265 at latencies as low as 8ms (4Kp60)
- 8K max resolution

#### AI-Enabled, Intelligent Video Pipeline

- Compression efficiency helps reduce bandwidth at optimal video quality
- AI-enabled “smart streaming” enhances quality of experience at lower bitrates

#### Cost-Effectively Scale for Interactive Media Applications

- Providers can affordably scale for high-volume, real-time streaming
- Enables a new class of interactive media and monetization services

As low as **8ms** Latency for 4K streaming (AV1)



### KEY APPLICATIONS

#### LIVE EVENTS

- Live Events
- Remote Production

#### VIDEO COLLABORATION

- Video Conferencing Streaming
- e-Learning
- Telemedicine
- Customer Service

#### CLOUD GAMING

- Cloud Gaming Streaming
- eSports

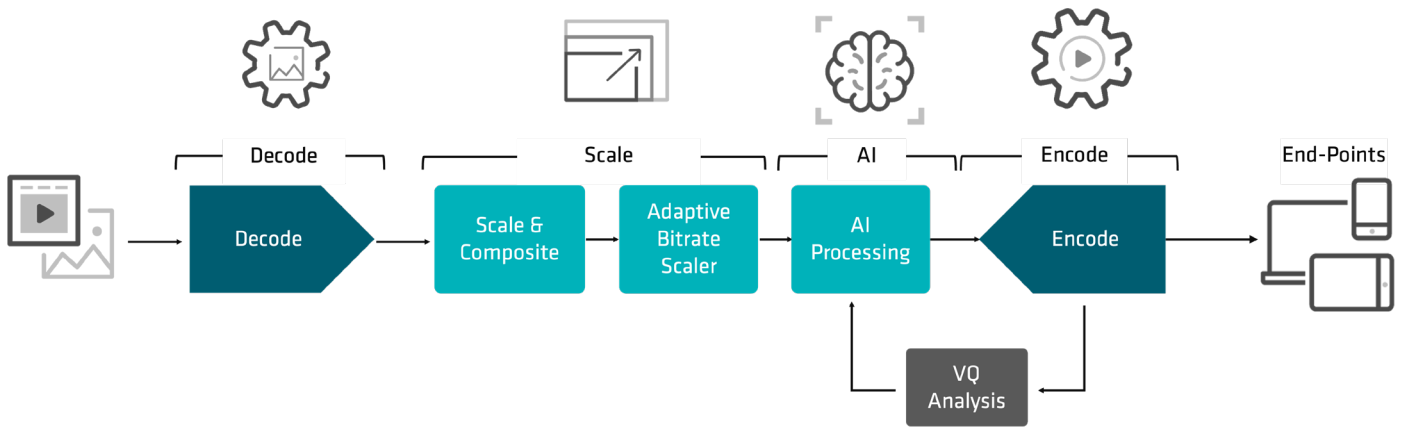
#### INTERACTIVE MEDIA

- Watch Parties
- Social Streaming (Content Creators)
- Live Shopping and eCommerce
- Online Auctions
- Online Wagering

## VIDEO PROCESSING UNIT FOR HIGH DENSITY, LOW LATENCY STREAMING

Designed from the ground-up for interactive media applications streaming, the Alveo™ MA350 media accelerator is based on a specialized ASIC architecture to deliver the best economics in terms of cost-, area-, and power-per-stream. Two of these purpose-built video processing units (VPU) power the accelerator and feature multiple video encoders, a hardened scaler to implement Adaptive Bitrate (ABR) ladders, a hardware compositor engine, an AI processor, video quality engines, and more. By accelerating the whole video pipeline on a single device, the Alveo MA350 platform reduces data movement between CPU and accelerator, maximizing channel density and minimizing chip-to-chip latency.

**Accelerating the Whole Video Pipeline**  
Video Processing Unit (2X per card)




The Alveo MA350 accelerator supports the stringent latency requirements of media applications and emerging use cases that offer personalization, gamification, and monetization services such as live shopping, live auctions, online wagering, and more. The card is architected such that latency reduces linearly with resolution and format reductions for diverse endpoints.



5sec ————— 8ms

### 4K STREAMING LATENCY CONTINUUM




## AI-ENABLED, INTELLIGENT VIDEO PIPELINE

The platform features an AI Processor per device that functions as a co-processor to the encoder. The dedicated engine evaluates content, frame-by-frame, and dynamically adjusts encoding (codec) parameters to enhance perceived visual quality while reducing bitrate. Optimization techniques include

- **Region-of-Interest (ROI)** encoding, where portions of the video are optimized—such as text and face—to enhance perceived visual quality while maintaining low bitrate.
- **Artifact Detection and Correction** removes ‘blockiness’, ringing, and blurring in video containing high level of motion or complexity.
- **Content Aware Encoding (CAE)** provides scene analysis predictive insights to encoder for bitrate optimization.

By integrating AI within a fully hardware accelerated pipeline, the Alveo MA35D can implement intelligent optimizations at scale and at low latency. The processor is inherently customizable to evolve for new AI techniques and use cases.



## COST-EFFECTIVELY SCALE FOR INTERACTIVE MEDIA APPLICATIONS

Profitably scaling high-volume streaming services requires maximizing the number of channels per server while minimizing power- and bandwidth-per stream. The Alveo MA35D media accelerator can deliver up to 32x 1080p streams per card, and with its slim Half-Height, Half-Length (HHHL) PCIe® form factor can equip a 1U rack server with 8-10 cards to maximize streams per server, per rack, and per data center. With 1 watt per stream<sup>1</sup> and 52% bitrate reduction compared to a software implementation<sup>2</sup>, the card helps reduce power and bandwidth OPEX so infrastructure and content providers can profitably scale.

### CAPEX

Up to 32x 1080p60 channels per card maximizes streams-per-server and helps reduce CAPEX




### Power OPEX

1 watt per stream<sup>1</sup> helps reduce overall power expenses



### Bandwidth OPEX

Up to 52% bitrate reduction<sup>2</sup> vs. software implementation for bandwidth savings



## SPECIFICATION

CODEC AND HARDWARE ENGINES	
<b>Encode Formats</b>	<ul style="list-style-type: none"> <li>• H.264   H.265   AV1</li> <li>• 2 discrete transcoding chips per card enable multiple standards concurrently (for new and legacy endpoints)</li> <li>• Faster than Real-Time encoding (file-based use case)</li> </ul>
<b>Decode Formats</b>	<ul style="list-style-type: none"> <li>• H.264   H.265   AV1   VP9</li> </ul>
<b>H.264 / H.265 Performance</b>	<ul style="list-style-type: none"> <li>• 4x 4Kp60   16x 1080p60   32x 1080p30   72x 720p30</li> </ul>
<b>AV1 Performance</b>	<ul style="list-style-type: none"> <li>• 4x 8Kp30   8x 4Kp60   32x 1080p60   64x 1080p30   144x 720p30</li> </ul>
<b>Max Channel Density &amp; Resolution</b>	<ul style="list-style-type: none"> <li>• 256 maximum transcode channels</li> <li>• 7680 x 4320 maximum resolution</li> </ul>
<b>8/10-bit Formats</b>	<ul style="list-style-type: none"> <li>• YUV20 planar and semi-planar</li> <li>• RGB planar</li> </ul>
<b>Comparable Presets</b>	<ul style="list-style-type: none"> <li>• H.264: x264 medium</li> <li>• H.265: x265 medium</li> <li>• AV1: x265 slow</li> </ul>
<b>Video Processing</b>	<ul style="list-style-type: none"> <li>• ABR Scaler</li> <li>• Compositor Engine</li> <li>• Video Look-Ahead Engine</li> <li>• Video Quality (VQ) and Quality-of-Experience (QoE) Engine</li> <li>• 2D graphics overlay</li> </ul>
<b>AI Processor</b>	<ul style="list-style-type: none"> <li>• AI Processor</li> <li>• 22 TOPS per card</li> <li>• Maximizes video quality at reduced bandwidth</li> <li>• AI-enabled Region-of-Interest optimization (e.g., text, face), artifact reduction, and content aware encoding</li> </ul>
<b>Auxiliary Processor</b>	<ul style="list-style-type: none"> <li>• 2x 64-bit quad-core RISC-V subsystem</li> <li>• Performs control and board management tasks</li> <li>• Minimizes software running on x86 host CPU</li> </ul>
BOARD SPECIFICATIONS	
<b>Form-Factor</b>	<ul style="list-style-type: none"> <li>• HHHH, Single-slot</li> </ul>
<b>Host Interface</b>	<ul style="list-style-type: none"> <li>• PCIe® Gen4 x8; bifurcated x4x4</li> <li>• SR-IOV</li> </ul>
<b>Memory</b>	<ul style="list-style-type: none"> <li>• 16GB LPDDR5</li> </ul>
<b>Typical Power</b>	<ul style="list-style-type: none"> <li>• 35W (50W TDP), passive cooling</li> </ul>

## TAKE THE NEXT STEP

- The Alveo MA35D media accelerator is now in production. Visit [amd.com/ma35d-inq](https://amd.com/ma35d-inq) to submit an inquiry or contact your sales representative.
- Download the new AMD Advanced Media Acceleration (AMA) SDK 1.0 for the MA35D available on [GitHub](https://github.com).



<sup>1</sup>Typical power for 8 4K streams or 32 1080p60 streams estimated at 35W, based on preliminary testing and subject to change. 50W Total Thermal Design Power (TDP)

<sup>2</sup>An Alveo™ MA35D AV1 Encode is 52% better on average in video quality than an open source x264 VeryFast SW model using various VMAF BD rates across 13 video files at various resolutions. Based on testing by AMD Labs in March 2023, using the VMAF scores of Alveo MA35D H.264 encode, H.265 encode, and AV1 encode compared to the VMAF scores of open source x264 VeryFast, X265 Medium, and X265 Slow SW models across (13) publicly available video files at various resolutions and bitrates. Actual results may vary. ALV-006

Video Multimethod Assessment Fusion (VMAF) is an objective full-reference video quality metric developed by Netflix in cooperation with the University of Southern California, predicting subjective video quality based on a reference and distorted video sequence. The metric can be used to evaluate the quality of different video codecs, encoders, encoding settings, or transmission variants.

## 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 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.

## COPYRIGHT NOTICE

© Copyright 2023 Advanced Micro Devices, Inc. All rights reserved. Xilinx, the Xilinx logo, AMD, the AMD Arrow logo, Alveo, Artix, Kintex, Kria, Spartan, Versal, Vitis, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. AMBA, AMBA Designer, ARM, ARM1176JZ-S, CoreSight, Cortex, and PrimeCell are trademarks of ARM in the EU and other countries. PCIe, and PCI Express are trademarks of PCI-SIG and used under license. PID2271900