

Overview

D/AVE NX is the latest and most powerful addition to the D/AVE family of rendering cores. It is the first IP to bring 3D graphics OpenGL ES 2.0 rendering (with some ES 3.0 / 3.1 extensions) to the FPGA and SoC world and – with offline-shader compilers – even into MCUs or low-end MPUs with small amounts of memory and bare-metal or RTOS operation systems.



Targeted for graphics applications on displays up to 4K x 4K resolution in the Industrial, Medical, Military, Avionics, Automotive and Consumer markets, **D/AVE NX** is designed to meet the sweet spot of performance and footprint bringing full 3D graphics even down to MCU class devices and into safety critical applications.

By enabling the use of programmable shaders even on small devices, high quality 2D and full 3D applications can be realized using the **D/AVE NX** core. Support for industry standard APIs like OpenGL ES 2.0 allows for rapid development of high-end user interfaces by leveraging common GUI frameworks like LVGL, Qt, SCADE or Guiliani and makes new, future-proof implementations possible.

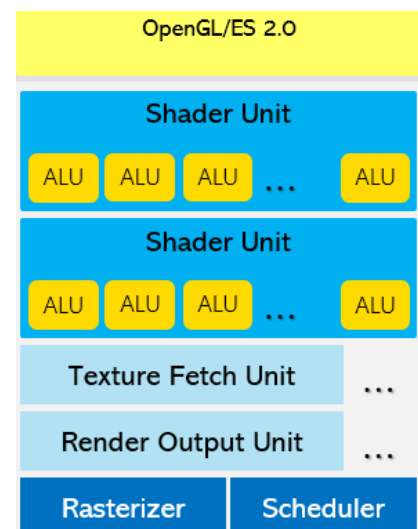
D/AVE NX can scale easily to fit exactly into the resource / performance sweet spot for a particular application. Entire device families can be equipped with differently scaled variants of the core, making all of them fully software-compatible. A single unified software stack and the guarantee to produce exactly the same visual result (at different speeds) allows saving significant development resources.

D/AVE NX is highly efficient as the internal multi-level scheduler can maximize the utilization of every hardware element even better than the fixed function pipeline of the successful D/AVE-2D and -HD cores could. Scheduling also does not have to be precomputed in the compiler, simplifying the compiler and driver architecture considerably.

Technology Details

System Features

- Scalability throughout the entire design
 - 1 - 4 shader units
 - 4 - 32 ALUs per shader unit (2 operations per ALU per cycle)
 - Optional Special-Function & I/O Units
 - Configurable cache sizes and optional 2nd level cache
 - Up to 4k resolution
 - Exact same driver can be used on all versions – enables the same output at different speeds
- Unified Shader Architecture
 - Dynamic, fully reconfigurable shaders
 - Efficient support for branches / divergent control flow
 - Fully IEEE compatible floating point ALUs
 - True integer arithmetic (8bit, 16bit, 32bit)
- Massively parallel execution with fine grained Multithreading
 - Thread context switch near instantaneous (~2 cycles)
 - Advanced task scheduler supporting both long- and short-term stalls
 - Architecture is able to eliminate hundreds of cycles of latency



- Concurrent use from multiple applications
- Job preemption possible on fragment level
- Immediate mode rendering (no need to buffer entire scene description)
- Application optimization and debugging support
 - Multiple hardware performance counters
 - Detailed analysis of shader stalls and scheduling
- Pipelined architecture for high clock frequencies
- System security features
 - Stop on bus error for integration with memory protection units
 - Hardware out-of-framebuffer memory access protection

Rendering

- OpenGL ES 2.0 API
 - with some OpenGL ES 3.0 / 3.1 extensions
 - Partial screen updates
 - OpenGL extensions for flash-based textures
 - Optional offline shader compilation
- High render quality
 - Highly accurate sub pixel positioning, interpolation, and filtering
 - Multiple anti-aliasing techniques (including MSAA)
- Texture Compression
- Framebuffer Compression
- Hardware supported blending
 - Normal alpha blending
 - Linear color-space blending
- Various texture and framebuffer formats
 - 8 bit alpha/luminance, ARGB4444, ARGB1555, RGB565, ARGB8888 etc.
 - Floating point texture support
 - 3d texture and texture array support
- Frame buffers and textures up to 4k x 4k pixels

Power Management

- Memory blocks controlled by Chip Select port
- Prepared for efficient automatic clock gating
- Global clock gating as option

Integration

- Single clock domain architecture
 - Bus interface clock frequency may differ from core frequency
- High latency capable
- Optional internal arbitration to work with a single bus master
- Adaptors for common bus protocols
 - ARM AMBA: APB for register access, AXI for memory bus master access
 - Altera Avalon as bus adaptors for both register and bus master access
 - Other bus protocols can be easily adapted

Performance and Resource Usage

The performance and resource usage of D/AVE NX depends mainly on the configuration concerning the number of Shader Units (SUs), the number of Arithmetic Logic Units (ALUs) per shader unit and partly on the bus and cache configuration.

Minimal footprint configuration

- 1 shader unit, 4 alus, 1k texture L1 cache, 512x512 resolution
- 4 GFLOPS at 500 MHz
- ~ 680 kGates(*) for Logic and 1393548 internal memory bits

Medium configuration

- 1 shader unit, 16 alus, 4k texture L1 cache, 1024x1024 resolution
- 16 GFLOPS at 500 MHz
- ~ 1215 kGates(*) for Logic and 2074092 internal memory bits

(*) 2-input NAND gate estimates derived from synthesis of D/AVE NX @600MHz with <10nm process

Verification Concept

A 100% algorithmic equivalent C model is used as reference for the verification of the RTL code. This real-time capable reference model is called 'SoftDAVE' which acts as a pixel accurate emulator on Windows PC. The emulator is also available as cost-free evaluation system and can be used for driver and application development.

Software

Drivers

TES provides Khronos conformant OpenGL ES 2.0 and EGL drivers. Both drivers rely on a low-level D/AVE NX driver layer, abstracting hardware details like register access and making porting to different CPUs / Operating systems easy.

All drivers have the following features:

- Fully reentrant & thread-safe
- Minimal OS dependency (HAL part separated)
- No inline assembler required
- Support for multiple D/AVE NX instances
- Multi-threading support, i.e. multiple applications can use D/AVE NX concurrently
- Small memory footprint

Memory footprint example - Stanford Bunny

- Example application running at 400 x 300
- Compressed framebuffer (stored at 8 bit per pixel)
- Compressed textures (stored at 4 bit per pixel)
- Single buffered (assuming display with internal memory)
- Full depth buffer (required only for true 3d applications)
- Complex shader to compute lighting and fine surface details



Application RAM


- 351 kB Frame/Depthbuffer
- 128 kB Textures

Driver RAM

- 32 kB Internal Buffers

Shader Compiler RAM


- 2520 kB Frontend
- 61 kB Backend



D/AVE NX reference Qt system solution on Yocto Linux

TES delivers a complete reference system solution for Altera SoCs (e.g. Arria 10 and Agilex 5) supporting selected reference boards including:

- Qt 5.x
- Yocto Linux OS
- OpenGL ES 2.0 and EGL drivers for D/AVE NX and CDC DRM as source code
- Qt and native OpenGL ES 2.0 example applications as source code
- Build scripts to check out required repositories (Yocto Linux, Qt, meta layers, drivers, ...) and build the D/AVE NX demo SD card image as well as the complete SDK.
- D/AVE NX as Megacore IP block (QSys component)
- CDC-200NX Display Controller as Megacore IP block (QSys component)



The package includes everything to evaluate the IP and start your own integration and application project.

Sales & Marketing Contact

TES Electronic Solutions GmbH
Wandalenweg 20
20097 Hamburg
Germany
<mailto:graphics@tes-dst.com>
www.tes-dst.com