



DATA SHEET

NVIDIA Tegra X1 Series Processors

Maxwell GPU + ARM v8

Description

The NVIDIA® Tegra® X1 series SoC couples the latest NVIDIA Maxwell™ GPU architecture with an ARM® v8 CPU cluster to deliver high performance and power efficiency. The Maxwell GPU implements a number of architectural enhancements designed to extract maximum performance per watt consumed. Designed for use in power-limited environments, Tegra X1 processors enable industry-leading visual computing capabilities, 32-bit and 64-bit operating capability, and integrate advanced multi-function audio, video and image processing pipelines.

Tegra X1 Series Processor*		
Description	TM670D	
	UCM: 1	UCM: 2
Use Case Model (UCM)		
% Operating time per day up to maximum UCM specification. Remaining time is spent in a SLEEP or OFF state.	20%	100%
Operating Temperature Range (Tj)	-25°C – 105°C	-25°C – 105°C
Operating Lifetime	5 years	5 years
Maxwell GPU		
256-core GPU End-to-end lossless compression Tile Caching OpenGL® 4.5 OpenGL ES 3.1 Vulkan™ 1.0 DirectX® 12 CUDA® 7.0 with FP16 GPGPU Android™ Extension Pack		
OpenGL ES Shader Performance (up to)	GFLOPS (FP16)	1024 ⁰
1024 ⁰		
CPU Subsystem		
ARM® Cortex® -A57 MPCore (Quad-Core) Processor with NEON Technology. Operating Frequency per Core (up to)	1.73GHz [‡]	1.68GHz [‡]
L1 Cache: 48KB L1 instruction cache (I-cache) per core; 32KB L1 data cache (D-cache) per core L2 Unified Cache: 2MB		
Memory Subsystem		
Dual Channel Secure External Memory Access Using TrustZone Technology System MMU		
Memory Type	4ch x 16-bit LPDDR4	
Maximum Memory Bus Frequency (up to) ^{††}	1600MHz	
Maximum Capacity	4GB	
HD Video & JPEG		
Decode [†]	H.265: Main10 H.265: Main H.264: Baseline, Main, High, Stereo SEI (half-res) H.264: MVC Stereo (per view) WEBM VP9 WEBM VP8 VC-1: Simple, Main, Advanced MPEG-2: Main	
Encode [†]	2160p 60fps 1080p 240fps 2160p 60fps 1080p 240fps 2160p 60fps 1080p 240fps 2160p 30fps 1080p 120fps 2160p 60fps 1080p 240fps 2160p 60fps 1080p 240fps 1080p 120fps 1080i 240fps 2160p 60fps 1080p 240fps 1080i 240fps	
JPEG (Decode & Encode)	600 MP/sec	
Audio Subsystem		
Dedicated programmable audio processor ARM Cortex A9 with NEON operating at up to 844MHz 10-input/5-output 8-channel audio mixer 4x Synchronous Sample Rate Converters		
Display Controller Subsystem		
Two independent display controllers with support for DSI with VESA link compression (VESA DSC), HDMI, DP and eDP		
Captive Panel		
MIPI-DSI (1.5Gbps/lane)	Uncompressed: 24bpp VESA DSC Compression: 12bpp	Support for Single x4 or Dual x4 links
eDP 1.4 (HBR2 5.4Gbps)	24bpp	Single link (1x4) 4096x2160 at 60Hz
External Display		
HDMI 2.0 (6Gbps)	24bpp	4096x2160 at 60Hz
DP 1.2a (HBR2 5.4Gbps)	24bpp	4096x2160 at 60Hz



Description	Tegra X1 Series Processor*	
	TM670D	
	UCM: 1	UCM: 2
Imaging System		
Dedicated RAW to YUV processing engines process up to 1400Mpix/s supports up to 24MP sensor		
MIPI CSI 2.0 up to 1.5Gbps (per lane)	Support for x4 and x2 configurations (up to 3 x4-lane or 6 x2-lane cameras)	
Clocks		
System clock: 38.4 MHz Sleep clock: 32.768 KHz Dynamic clock scaling and clock source selection		
Boot Sources		
eMMC: IROM (Primary) and USB (Recovery Mode)		
Security		
Secure memory with video protection region for protection of intermediate results Configurable secure DRAM regions for code and data protection Hardware acceleration for AES 128/192/256 encryption and decryption to be used for secure boot and multimedia Digital Rights Management (DRM) Hardware acceleration for AES CMAC, SHA-1 and SHA-256 algorithms 2048-bit RSA HW for PKC boot HW Random number generator (RNG) SP800-90 TrustZone technology support for DRAM, peripherals Dedicated HDCP HW		
Storage		
4 x SD/MMC controllers (supporting eMMC 5.1, SD 4.0, SDHOST 4.0 and SDIO 3.0) SATA		
Peripheral Interfaces*Δ		
xHCI host controller with integrated PHY: 3 x USB 3.0, 3 x USB 2.0 USB 3.0 device controller with integrated PHY EHCI controller with embedded hub for USB2.0 5-lane PCIe: one x1 and one x4 controllers SATA (1 port) 4 x UART 3 x SPI 6 x I²C 4 x I2S: support I²S, RJM, LJM, PCM, TDM (multi-slot mode)		
Applications		
Embedded (Intelligent Video Analytics, Drones, Robotics, etc.), Automotive, Clamshells, Gaming, Internet TV, and more		

* Refer to the software release feature list for current software support; not all features are exposed on all OS options.

◊ See Table 2 for Guaranteed GPU operation across supported temperature range.

‡ See Table 1 for Guaranteed CPU operating frequency across supported temperature range.

†† Dependent on board layout, refer to Interface Design Guide for layout guidelines.

Δ USB3.0, PCIe, and SATA availability are determined on platform configuration due to PHY lane sharing. Refer to the Interface Design Guide and Technical Reference Manual to determine which peripheral interface options can be simultaneously exposed.



Revision History

Version	Date	Description
v0.99	NOV, 2015	Initial Release.
v1.0	FEB, 2016	CPU Complex: updated Guaranteed CPU Operating Frequency.

Table of Contents

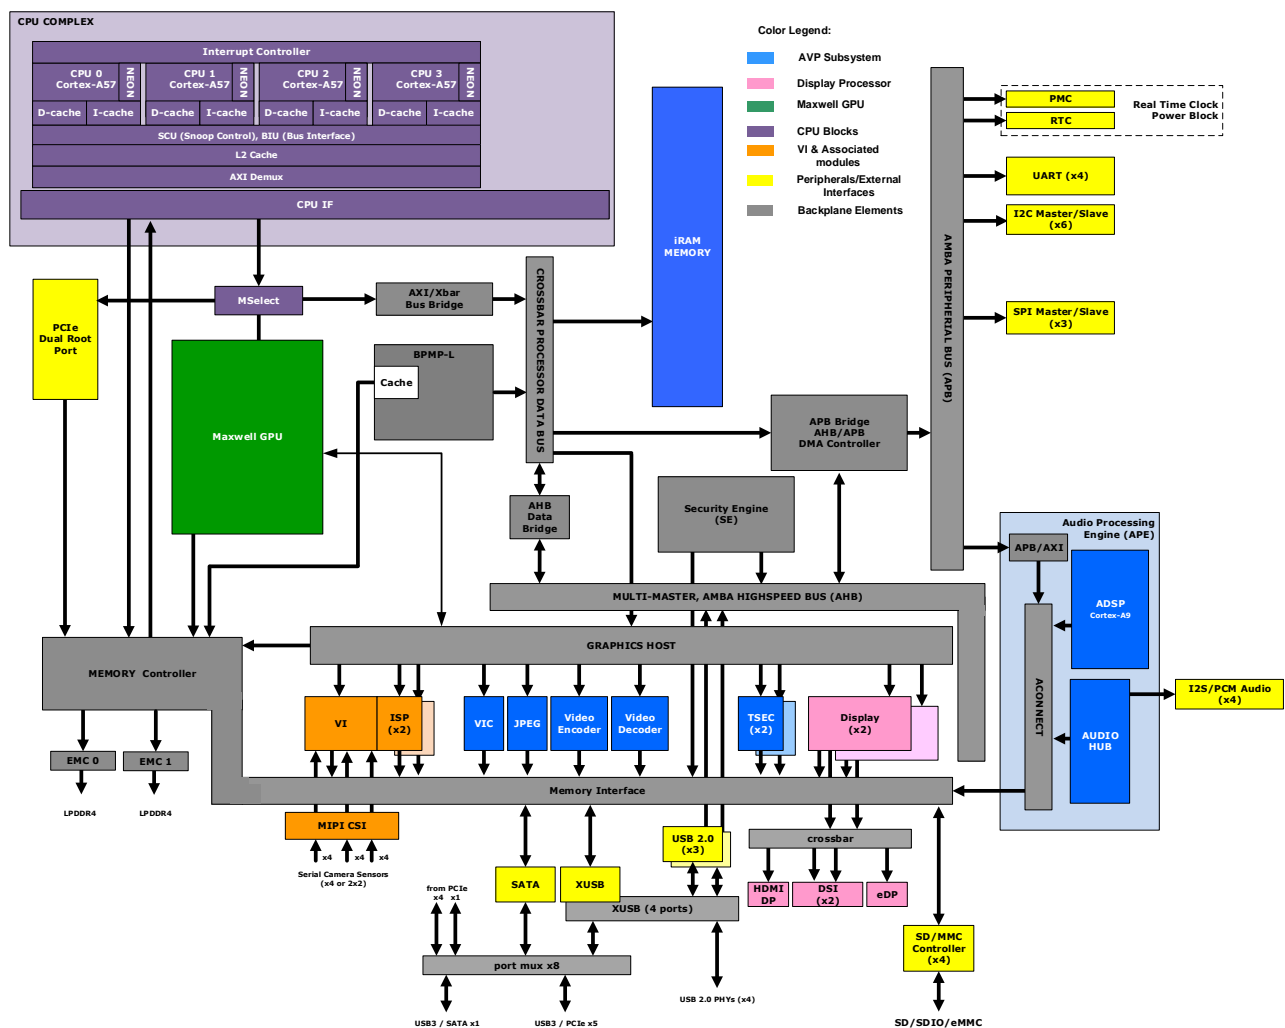
1.0 Tegra X1 SoC Overview	6
1.1 CPU Complex	7
1.1.1 Snoop Control Unit and L2 Cache	7
1.1.2 Performance Monitoring	8
1.2 High-Definition Audio-Video Subsystem	8
1.2.1 Multi-Standard Video Decoder	8
1.2.2 Multi-Standard Video Encoder	8
1.2.3 JPEG Processing Block	9
1.2.4 Video Image Compositor (VIC)	9
1.2.5 Audio Processing Engine (APE)	10
1.2.6 Tegra Security Controller (TSEC)	10
1.3 Maxwell GPU	11
1.4 Image Signal Processor (ISP)	12
1.5 Display Controller Complex	13
1.6 Memory Controller	14
1.7 Security Engine	14
2.0 Power and System Management	15
2.1 Power Domains/Islands	15
2.2 Power Management Controller (PMC)	15
2.2.1 Resets	15
2.2.2 System Power States and Transitions	16
2.3 Thermal & Power Monitoring	17
3.0 Interface and Signal Descriptions	18
3.1 External Memory Controller (EMC)	18
3.2 SD/eMMC Controller	18
3.3 Serial ATA (SATA) Controller	19
3.4 Display Interfaces	19
3.4.1 MIPI Display Serial Interface (DSI)	19
3.4.2 High-Definition Multimedia Interface (HDMI) and DisplayPort (DP) Interfaces	20
3.4.3 Embedded DisplayPort (eDP) Interface	21
3.5 Audio Interfaces: I2S, PCM, TDM	21
3.6 USB Interfaces	22
3.7 PCI Express (PCIe) Interface	22
3.8 Serial Peripheral Interface	22
3.9 Inter-Chip Communication (I2C) Controller	24
3.10 UART Controller	24
3.11 Video Input Interfaces	25
3.11.1 MIPI Camera Serial Interface (CSI)	25
3.11.2 Camera / VI (Video Input)	26
3.12 Debug Interface	26

The Tegra X1 series processor integrates an NVIDIA Maxwell GPU coupled to quad-core ARM Cortex-A57 CPU. This section provides a brief overview of the processing blocks listed below:

- CPU Complex
- High-Definition Audio-Video Subsystem
- Maxwell GPU
- Image Signal Processor (ISP)
- Display Controller Complex
- Security Controller
- Memory Controller

Descriptions of peripheral interfaces exposed on the Jetson TX1 module can be found in the *Interface Descriptions* section.

Figure 1 Tegra X1 SoC Block Diagram



1.1 CPU Complex

The CPU complex is a high performance Multi-Core SMP cluster of four Cortex-A57 CPUs with 2MB of L2 cache (shared by all cores). Features include:

- Superscalar, variable-length, out-of-order pipeline
- Dynamic branch prediction with Branch Target Buffer (BTB) and Global History Buffer RAMs, a return stack, and an indirect predictor
- 48-entry fully-associative L1 instruction TLB with native support for 4KB, 64KB, and 1MB page sizes.
- 32-entry fully-associative L1 data TLB with native support for 4KB, 64KB and 1MB pages sizes.
- 4-way set-associative unified 1024-entry Level 2 (L2) TLB in each processor
- 48Kbyte I-cache and 32Kbyte D-cache for each core.
- Full implementation of ARMv8 architecture instruction set
- Embedded Trace Microcell (ETM) based on the ETMv4 architecture
- Performance Monitor Unit (PMU) based on the PMUv3 architecture
- Cross Trigger Interface (CTI) for multiprocessor debugging
- Cryptographic Engine for crypto function support
- Interface to an external Generic Interrupt Controller (vGIC-400)
- Power management with multiple power domains

The CPU cluster interface to the MSelect FIFO via an AXI interface to decouple I/O traffic. MSelect allows an AXI master device to send traffic to the peripheral buses based on transaction address. The AXI/Xbar bridge enables early response on write transfers and full hardware hazard resolution to permit the maximum transaction throughput to MMIO.

Table 1 lists the CPU frequencies that the Tegra X1 processor is capable of supporting at different temperatures. CPU frequency and voltage are actively managed by Tegra Power and Thermal Management Software and influenced by workload. Frequency may be throttled at higher temperatures (over 70C) resulting in a behavior that reduces the CPU operating frequency. Observed chip-to-chip variance is due to NVIDIA ability to maximize performance (DVFS) on a per-chip basis, within the available power budget.

Table 1 TM670D (UCM #1, UCM #2) Guaranteed CPU Operating Frequency

Use Case Model		Guaranteed CPU Operating Frequency		
UCM	V _{MIN} (V)	T _j ≤ 70C	T _j ≤ 90C	T _j ≤ 105C
UCM: 1 20% operating time per day up to maximum specification. Remaining time is spent in a SLEEP or OFF state.	0.841	1.73GHz	1.63GHz	1.55GHz
UCM: 2 100% operating time per day up to maximum specification.	0.90	1.68GHz	1.55GHz	1.47GHz

1.1.1 Snoop Control Unit and L2 Cache

The CPU cluster includes an integrated snoop control unit (SCU) that maintains coherency between the CPUs within the cluster and a tightly coupled L2 cache that is shared between the CPUs within the cluster. The L2 cache also provides a 128-bit AXI master interface to access DRAM. L2 cache features include:

- 2MB L2
- Fixed line length of 64 bytes
- 16-way set-associative cache structure
- Duplicate copies of the L1 data cache directories for coherency support
- Hardware pre-fetch support
- ECC support

1.1.2 Performance Monitoring

The performance monitoring unit (part of MPCore non-CPU logic) provides six counters, each of which can count any of the events in the processor. The unit gathers various statistics on the operation of the processor and memory system during runtime, based on ARM PMUv3 architecture.

1.2 High-Definition Audio-Video Subsystem

The audio-video subsystem off-loads audio and video processing activities from the CPU subsystem resulting in faster, fully concurrent, highly efficient operation.

1.2.1 Multi-Standard Video Decoder

The video decoder accelerates video decode, supporting low resolution content, Standard Definition (SD), High Definition (HD) and UltraHD (2160p, or “4k” video) profiles. The video decoder is designed to be extremely power efficient without sacrificing performance.

The video decoder communicates with the memory controller through the video DMA which supports a variety of memory format output options. For low power operations, the video decoder can operate at the lowest possible frequency while maintaining real-time decoding using dynamic frequency scaling techniques.

Video standards supported:

- Encrypted content for: H.264, H.265, VP9
- H.265: Main10, Main
- WEBM VP9 and VP8
- H.264: Baseline (no FMO/ASO support), Main, High, Stereo SEI (half-res)
- VC-1: Simple, Main, Advanced
- MPEG-4: Simple (with B frames, interlaced; no DP and RVLC)
- H.263: Profile 0
- DivX: 4 / 5 / 6
- XviD Home Theater
- MPEG-2: MP

1.2.2 Multi-Standard Video Encoder

The multi-standard video encoder enables full hardware acceleration of various encoding standards. It performs high-quality video encoding operations for applications such as video recording and video conferencing. The encode processor is designed to be extremely power-efficient without sacrificing performance.

Video standards supported:

- H.265 Main Profile: I-frames and P-frames (No B-frames)
- H.264 Baseline/Main/High Profiles: IDR/I/P/B-frame support, MVC
- VP8
- MPEG4 (ME only)
- MPEG2 (ME only)
- VC1 (ME only): No B frame, no interlaced

1.2.3 JPEG Processing Block

The JPEG processing block is responsible for JPEG (de)compression calculations (based on JPEG still image standard), image scaling, decoding (YUV420, YUV422H/V, YUV444, YUV400) and color space conversion (RGB to YUV; decode only).

Input (encode) formats:

- Pixel width: 8bpc
- Subsample format: YUV420
- Resolution up to 16K x 16K
- Pixel pack format
 - Semi-planar/planar for 420

Output (decode) Formats:

- Pixel width 8bpc
- Resolution up to 16K x 16K
- Pixel pack format
 - Semi-planar/planar for YUV420
 - YUY2/planar for 422H/422V
 - Planar for YUV444
 - Interleave for RGBA

1.2.4 Video Image Compositor (VIC)

The Video Image Compositor implements various 2D image and video operations in a power-efficient manner. It handles various system UI scaling, blending and rotation operations, video post-processing functions needed during video playback, and advanced de-noising functions used for camera capture.

Features:

- Color Decompression
- High-quality Deinterlacing
- Inverse Teleciné
- Temporal Noise Reduction
 - High quality video playback
 - Reduces camera sensor noise
- Scaling
- Color Conversion
- Memory Format Conversion
- Blend/Composite
- 2D Bit BLIT operation
- Rotation

1.2.5 Audio Processing Engine (APE)

The Audio Processing Engine (APE) is a self-contained unit that provides a complete audio solution. The APE includes the Audio Digital Signal Processor (ADSP), Audio Hub (AHUB) and Audio Connect (ACONNECT). Software based post processing effects enable the ability to implement custom audio algorithms.

Features:

- Audio Digital Signal Processor (ADSP)
 - ARM Cortex-A9
 - NEON SIMD & FPU
 - 32K-I/32K-D L1, 128K L2 cache
- 64KB Audio RAM
- Dedicated audio clocking enables ULP audio processing
- Low latency voice processing
- Audio Hub (AHUB)
 - 3 x I2S Stereo I/O
 - PDM Receiver: 3 x (Stereo) or 6 x (Mono)
- Multi-Channel IN/OUT
 - Digital Audio Mixer: 10-in/5-out
- Up to 8 channels per stream
- Simultaneous Multi-streams
- Flexible stream routing
 - Built-in speaker protection with I/V sensing
 - Multi-band Dynamic Range Compression (DRC)
- Up to 3 bands
- Customizable DRC curve with tunable knee points
- Up to 192KHz, 32-bit sample, 8 channels
 - Parametric equalizer: up to 12 bands
 - Low latency sample rate conversion (SRC)

1.2.6 Tegra Security Controller (TSEC)

Tegra Security Controllers (TSEC) are dedicated hardware accelerators used to support secure content transfer and playback. TSEC heavy-secure (HS) hardware is capable of authenticating its own code autonomously using its Secure Boot ROM and signature verification keys. The on-chip secure memory enables tamper resistant secure storage and transaction verification. TSEC implements a random number generator (RNG) and AES-128 encryption/decryption; no other cryptographic primitives or key sizes are supported. Two independent instruction queues (capable of holding up to 16 instructions) are used to provide encryption support for DRM schemes, including protected content encryption/ decryption and link management.

Tegra X1 processors incorporate two instances of the TSEC controller to balance the performance requirements of increasingly demanding use cases.

- TSECA
 - HDCP 2.2 over HDMI 2.0 (wired)
 - HDCP 1.4 over HDMI 1.0 (wired)
- TSECB
 - NVIDIA Security Interface (NVSI)

Features:

- HDCP Link Management
 - HDCP link management without exposing protected content or HDCP keys to SW running on CPU.
 - (Programmable) Ability to disable HDMI output independent of the player if the HDCP status check fails.
- WiFi Alliance Display (WFD) Encryption (TSECB Only)
 - Requires HDCP 2.0 encryption support
 - Ability to maintain 30 FPS video rate in video pipe
- NVIDIA Security Interface (TSECB Only): NVSI is a gaming DRM scheme supported in Tegra products that leverages the Google® Play Store for content download. NVSI-support is included royalty-free in Tegra products and has been designed to prevent game titles from being copied.
- Dedicated Video Protection Region in memory (TSECB Only)
 - Programmable in the memory controller
 - Extends security controller i-cache and d-cache
 - Only accessible by the Security Controller
 - Minimum size requirements avoid security exposure
- Blu-Ray/MPEG2-TS playback
 - Decrypt and parse Blu-Ray/MPEG2-TS streams
 - Encrypt video stream using AES and write the encrypted stream to memory.
 - Read/write to the Video Protection Region

1.3 Maxwell GPU

The Maxwell GPU introduces an all-new design for the Streaming Multiprocessor (SM) that dramatically improves performance per watt and performance per area. On the surface, Maxwell GPUs appear similar to the previous Kepler generation; the same APIs are used, there are multiple SM units within a Graphics Processing Cluster (GPC), each SM includes a Polymorph Engine and Texture Units, each GPC includes a Raster Engine. ROPs are still aligned with L2 cache slices and Memory Controllers. Internally, the Maxwell GPU redesigned all unit and crossbar structures, optimized data flows, and significantly improved power management. The SM scheduler architecture and algorithms were rewritten to be more intelligent and avoid unnecessary stalls, while further reducing the energy per instruction required for scheduling. The organization of the SM also changed. Each Maxwell SM (called SMM) is now partitioned into four separate processing blocks, each with its own instruction buffer, scheduler and 32 CUDA cores.

The GPC is a dedicated hardware block for rasterization, shading, texturing, and compute; most of the GPU's core graphics functions are performed inside the GPC. Inside the GPC, the SMM CUDA cores perform pixel/vertex/geometry shading and physics/compute calculations. Texture units perform texture filtering and load/store units fetch and save data to memory. Special Function Units (SFUs) handle transcendental and graphics interpolation instructions. Finally, the PolyMorph Engine handles vertex fetch, tessellation, viewport transform, attribute setup, and stream output. The SMM geometry and pixel processing performance make it highly suitable for rendering advanced user interfaces and complex gaming applications; the power efficiency of the Maxwell GPU enables this performance on devices with power-limited environments.

Features:

- End-to-end lossless compression
- Tile Caching
- Support for OpenGL 4.5, OpenGL ES 3.1 + AEP (Android Extension Pack), Vulkan 1.0, DirectX 12, CUDA 7.0 (FP16)
- Adaptive Scalable Texture Compression (ATSC) LDR profile supported

- Iterated blend, ROP OpenGL-ES blend modes
- 2D BLIT from 3D class avoids channel switch
- 2D color compression
- Constant color render SM bypass
- 2x, 4x, 8x MSAA with color and Z compression
- Non-power-of-2 and 3D textures, FP16 texture filtering
- FP16 shader support
- Geometry and Vertex attribute Instancing
- Parallel pixel processing
- Early-z reject: Fast rejection of occluded pixels acts as multiplier on pixel shader and texture performance while saving power and bandwidth
- Video protection region
- Power saving: Multiple levels of clock gating for linear scaling of power

Table 2 lists the GFLOP operations that the Tegra X1 processor is capable of supporting at different temperatures. GPU frequency and voltage are actively managed by Tegra Power and Thermal Management Software and influenced by workload. Frequency may be throttled at higher temperatures (over 70C) resulting in a behavior that reduces the GPU operating frequency. Observed chip-to-chip variance is due to NVIDIA ability to maximize performance (DVFS) on a per-chip basis, within the available power budget.

Table 2 TM670D (UCM#1, UCM#2) Guaranteed GPU Operation

Use Case Model	Guaranteed GPU Operations: GFLOPS (FP16)		
	Tj <= 70C	Tj <= 90C	Tj <= 105C
UCM: 1 20% operating time per day up to maximum specification. Remaining time is spent in a SLEEP or OFF state.	1024	1024	942
UCM: 2 100% operating time per day up to maximum specification.	1024	1024	942

1.4 Image Signal Processor (ISP)

The ISP module takes data from the VI/CSI module or memory in raw Bayer format and processes it to YUV output. The imaging subsystem supports raw (Bayer) image sensors up to 24 million pixels. Advanced image processing is used to convert input to YUV data and remove artifacts introduced by high megapixel CMOS sensors and optics with up to 30-degree CRA.

Features:

- Flexible post-processing architecture for supporting custom computer vision and computational imaging operations
- Bayer domain hardware noise reduction
- Per-channel black-level compensation
- High-order lens-shading compensation
- 3x3 color transform
- Bad pixel correction
- Programmable coefficients for de-mosaic with color artifact reduction
 - Color Artifact Reduction: a two-level (horizontal and vertical) low-pass filtering scheme that is used to reduce/remove any color artifacts that may result from Bayer signal processing and the effects of sampling an image.
- Enhanced down scaling quality
- Edge Enhancement

- Color and gamma correction
- Programmable transfer function curve
- Color-space conversion (RGB to YUV)
- Image statistics gathering (per-channel)
 - Two 256-bin image histograms
 - Up to 4,096 local region averages
 - AC flicker detection (50 Hz and 60 Hz)
 - Focus metric block

1.5 Display Controller Complex

The Tegra Display Controller Complex integrates two independent display controllers. Each display controller is capable of interfacing to an external display device and can drive the same or different display contents at different resolutions and refresh rates. Each controller supports a cursor and three windows (Window A, B, and C); controller A supports two additional simple windows (Window D,T). The display controller reads rendered graphics or video frame buffers in memory, blends them and sends them to the display.

Features:

- Two heads. Each can be mapped to one of:
 - SOR0: (no audio/HDCP) HDMI/DP/eDP
 - SOR1: HDMI/DP/eDP
 - DSI A/B (up to 8-lanes with ganged mode)
- 90, 180, 270-degree image transformation uses both horizontal and vertical flips (controller A only)
- Byte-swapping options on 16-bit and 32-bit boundary for all color depths
- NVIDIA Pixel Rendering Intensity and Saturation Management™ (PRISM)
- 256x256 cursor size
- Color Management Unit for color decompression and to enhance color accuracy (compensate for the color error specific to the display panel being used)
- Scaling and tiling in HW for lower power operation
- Full color alpha-blending
- Captive panels
 - Secure window (Win T) for TrustZone
 - Supports cursor and up to four windows (Win A, B, C and D)
 - 2x4-lane MIPI DSI (supports a single Hi-Res panel in 2x4 ganged mode). 2x4 can support left-right, odd-even split configurations.
 - Supports MIPI D-PHY rates up to 1.5Gbps
 - 4-lane eDP with AUX channel
 - Independent resolution and pixel clock
 - Supports display rotation and scaling in HW
- External displays
 - Supports cursor and three windows (Window A, B, and C)
 - 1x HDMI (2.0) or DisplayPort (HBR2) interface
 - Supports display scaling in HW

1.6 Memory Controller

The Tegra Memory Controller (MC) maximizes memory utilization while providing minimum latency access for critical CPU requests. An arbiter is used to prioritize requests, optimizing memory access efficiency and utilization and minimizing system power consumption. The MC provides access to main memory for all internal devices. It provides an abstract view of memory to its clients via standardized interfaces, allowing the clients to ignore details of the memory hierarchy. It optimizes access to shared memory resources, balancing latency and efficiency to provide best system performance, based on programmable parameters.

Features:

- TrustZone (TZ) Secure and OS-protection regions
- System Memory Management Unit
- Dual CKE signals for dynamic power down per device
- Support for two DRAM ranks of unequal device densities
- Dynamic Entry/Exit from Self -Refresh and Power Down states

The Tegra MC is able to sustain high utilization over a very diverse mix of requests. For example, the MC is prioritized for bandwidth (BW) over latency for all multimedia blocks (the multimedia blocks have been architected to prefetch and pipeline their operations to increase latency tolerance); this enables the MC to optimize performance by coalescing, reordering, and grouping requests to minimize memory power. DRAM also has modes for saving power when it is either not being used, or during periods of specific types of use.

1.7 Security Engine

A dedicated platform security engine supports secure boot using AES or PKC, incorporates an SP800-90 compliant random number generator (RNG) including built in ring oscillator based entropy source used to seed a deterministic random bit generator (DRBG), and a protected memory aperture for video use cases.

Features:

- Streaming memory-to-memory and on-the-fly (OTF) AES decryption
 - Modes: ECB, CBC, OFB, CTR
 - Hash: CMAC
- Secure boot
 - AES: Boot configuration table (BCT) and Boot Loader (BL) are decrypted/authenticated/loaded into memory; Boot ROM locks down security features and clears out state; BL write protects mass storage location of BL and OS
 - PKC: Boot ROM performs 2048-bit RSA signature verification; once the public key stored in mass storage is validated, the key is used to verify the BCT/BL hash.
- Secure memory
 - Secure ROM: regions locked before control given to BL
 - Secure RAM (TZRAM): Security controlled by Secure OS (TZ) Tasks
 - Secure DRAM access
 - Video Protection Region: MC dynamically configures memory region that can only be accessed by video/display HW engines.
- AES key slot protection
 - Protection scheme associated with each key slot defines Read/Write permissions
 - Configure key slots so that they can only be accessed in TZ mode
 - Individual key slots can be set so they can only be used for AES operations by TZ processes
 - Valid AES decryption destination
- TZ Secured Peripheral Bus
- HW Hashing & Authentication: AES CMAC, SHA1, SHA2

2.0 Power and System Management

Tegra processors utilize various means to provide an efficient power management solution for a complex environment:

- **Power Management Controller (PMC) and Real Time Clock (RTC):** These blocks reside in an Always On (not power gated) partition. The PMC provides an interface to an external power manager IC or PMU. It primarily controls voltage transitions for the Tegra processor as it transitions to/from different low power modes; it also acts as a slave receiving dedicated power/clock request signals as well as wake events from various sources (e.g., SPI, I2C, RTC, USB attach) which can wake the Tegra processor from a deep sleep state. The RTC maintains the ability to wake the system based on either a timer event or an external trigger (e.g., key press).
- **Power Gating:** Tegra processor aggressively employs power-gating (controlled by PMC) to power-off modules which are idle. CPU cores are on a separate power rail to allow complete removal of power and eliminate leakage. Each CPU can be power gated independently. Software provides context save/restore to/from DRAM.
- **Clock Gating:** Used to reduce dynamic power in a variety of power states.
- **Dynamic Voltage and Frequency Scaling (DVFS):** Raises voltages and clock frequencies when demand requires, Lowers them when less is sufficient, and removes them when none is needed. DVFS is used to change the voltage and frequencies on the following rails: VDD_CPU, VDD_CORE and VDD_GPU.

2.1 Power Domains/Islands

Tegra processors are partitioned into power domains and power islands to optimize mobile device battery life by reducing power consumption for various use cases and limiting leakage current.

Tegra X1 series processors have four power domains (RTC/CORE/GPU/CPU); RTC domain is always on, CORE/CPU/GPU domains can be turned on and off. The CPU, CORE and GPU power domains contain power-gated islands which are used to power individual modules (as needed) within each domain. Clock-gating is additionally applied during powered-on but idle periods to further reduce unnecessary power consumption and such clock-gating is applied to both power-gated and non-power-gated islands (NPG).

2.2 Power Management Controller (PMC)

The PMC interacts with an external Power Manager IC or PMU through side band signals. It incorporates power management features that enable both high speed operation and very low-power standby states. The PMC primarily controls voltage transitions for the Tegra processor as it transitions to/from different low power modes; it also acts as a slave receiving dedicated power/clock request signals as well as wake events from various sources (e.g., SPI, I2C, RTC, USB attach) which can wake the Tegra processor from deep sleep state. Tegra processors (with PMC support) are able to employ aggressive power-gating capabilities on idle modules.

The PMC integrates specific logic to maintain defined states and control power domains (including signaling the external PMU to provide power) during sleep and deep sleep modes.

2.2.1 Resets

The PMC receives the primary chip reset event (from SYS_RESET_N) and generates various resets for: PMC, RTC and CAR. From the PMC provided reset, the Clock and Reset (CAR) controller generates resets for most of the blocks in the chip. In addition to chip reset events, the PMC receives other events (e.g., thermal, WatchDog Timer (WDT), SW, wake) which also result in variants of system reset.

The RTC block includes an embedded real-time clock and can wake the system based on either a timer event or an external trigger (e.g., key press).

2.2.2 System Power States and Transitions

Because of the aggressive power gating and monitoring techniques that individual blocks implement there are many different methods used to define when some blocks in the SoC are idle or powered.

Table 3 System Power States

Name	Description
ACTIVE (SC0)	System is running under DVFS control <ul style="list-style-type: none"> CPU, Devices, and System clocks are dynamically scaled Full functionality available
IDLE (SC1)	HW managed state The most shallow idle state. In this state memory is put into self-refresh mode, CPU power-rail is OFF. This is the only HW managed SC state; state through HW DSR mechanism (Dynamic Self Refresh). All deeper states are SW (or BPMP FW) based.
IDLE (SC2)	SW managed state The first SW (or BPMP FW) managed state. In this state memory clock network is turned off, prerequisite for entering this state is turning off all the DMA engines. SC state SW gets this information by monitoring Runtime Power Management (RPM) statuses of all the relevant drivers. <ul style="list-style-type: none"> Software controlled self-refresh is enabled CPU power-rail is OFF
IDLE (SC3)	SW managed state <ul style="list-style-type: none"> If current memory frequency does not require PLLM source SW/FW turns off PLLM. Reduces power of the ubiquitous PLLP clock network by switching branches that allow sudden reduction in frequency to OSC clock (which runs at lower frequency). CPU power-rail is OFF
SUSPEND (SC4)	SW managed state <ul style="list-style-type: none"> Inactivity timeout, no CPU process needed, no devices are active OS is suspended DRAM is in self-refresh CPU power-rail is OFF. SoC power-gate-able units (incl. CPULP) are power-gated GPU rail is OFF OSC clock is off: clock input is clamped (optional). Wake events (incl. interrupts) are possible Some Tegra devices may be available, but at extremely low performance
DEEP SLEEP (SC7)	Inactivity timeout, no CPU process needed, no devices are active. PADs are powered off except for PADs which monitor wake events <ul style="list-style-type: none"> OS is suspended DRAM is in self-refresh Both CPU and SoC power rails are OFF GPU rail is OFF OSC clock is off PMC and RTC still available. Wake events (incl. USB) are possible.
OFF	Tegra system (incl. DRAM) is completely powered off. <ul style="list-style-type: none"> No state is maintained. No internal wake events possible.

2.3 Thermal & Power Monitoring

Tegra X1 series processors employ both integrated and external thermal sensors to monitor temperature:

- Multiple temperature-sensitive oscillators (TSOSC) placed in different partitions on the SoC for on-chip thermal sensing.
- An always-on thermal alarm generator (AOTAG) monitors die temperature and compares that temperature to a low/high-temperature threshold and initiates action if the temperature exceeds the threshold.
- An on-chip thermal diode (TDIODE) enables ability to connect an external thermal sensor. External signals drive a constant current through the diode and measure the voltage drop, which is a function of temperature.

Tegra X1 series processors take actions based on the current die temperature:

- Temperature is within safe operation conditions (TH0): No action taken.
- Temperature exceeds threshold temperature 1 (TH1): An interrupt is generated for SW. SW DVFS dynamically adjusts frequency and voltage through SW controls to halt temperature rise. When temperature returns to below TH1, DVFS clock frequency reductions are reverted.
- Temperature exceeds threshold temperature 2 (TH2): The frequencies of the CPUs and GPU are immediately cut by a preprogrammed amount and an interrupt is sent to SW. When temperature returns to below TH2, clock frequency reductions are reverted.
- Temperature exceeds threshold temperature 3 (TH3): A signal is sent to PMC to reset the chip. A flag is set in the PMC to indicate the chip shut down due to a thermal event.

Tegra X1 series processors also support power monitoring functionality. The SOC_THERM block can be configured to throttle both the CPU and GPU in response to any over-current, under-voltage or over-power inputs from the platform.

3.0 Interface and Signal Descriptions

This section describes device signals. Additional alternate use signals are described in the *Tegra X1 Series Processor Technical Reference Manual*. Signals are arranged in functional groups according to their associated interface.

3.1 External Memory Controller (EMC)

Features:

- LPDDR4: supports x16 DRAM chips
- 64-bit data bus
- BL16 support
- Low Latency Path and Fast Read/Response Path Support for the CPU Complex Cluster
- Support for low-power modes:
 - Software controllable entry/exit from: self-refresh, power down, deep power down
 - Hardware dynamic entry/exit from: power down, self-refresh
 - Support for intermittent or disabled DLL
 - Disable unused address/command taps based on mode
 - Pads use DPD-mode during idle periods

3.2 SD/eMMC Controller

Standard	Notes
<i>SD Specifications Part A2 SD Host Controller Standard Specification Version 4.00</i>	
<i>SD Specifications Part 1 Physical Layer Specification Version 4.00</i>	
<i>SD Specifications Part E1 SDIO Specification Version 4.00</i>	Support for SD 4.0 Specification without UHS-II
<i>Embedded Multimedia Card (eMMC), Electrical Standard 5.01</i>	

The SecureDigital (SD)/Embedded MultiMediaCard (eMMC) controller is capable of interfacing to SD/eSD, SDIO cards, and eMMC devices. It has a direct memory interface and is capable of initiating data transfers between memory and external card. The SD/eMMC controller supports 2 different bus protocols: SD and eMMC bus protocol for eMMC. It has an APB Slave interface to access configuration registers. To access the iRAM for Micro Boot, the SD/eMMC controller relies on the AHB redirection arbiter in the Memory Controller.

Features:

- Supports of 8-bit data interface for eMMC/eSD cards
- Supports 4-bit data interface for SD cards
- Allows card to interrupt host in 1-bit, 4-bit, 8-bit SD modes.
- Supports Read wait Control, Suspend/Resume operation for SDIO cards
- Supports FIFO overrun and underrun condition by stopping SD clock
- Supports addressing larger capacity SD 3.0 or SD-XC cards up to 2 TB.

Tegra X1 series processors support four instances of this controller. These controllers can be routed to multiple physical locations on the device. The SD/SDIO controllers support Default and High Speed modes as well as the High and Low voltage ranges.

Table 4 SD/MMC Controller I/O Capabilities

Controller	Bus Width	eMMC 5.01 Support	Supported Voltages (V)	I/O bus clock (MHz)	Maximum Bandwidth (MBps)	Notes
SDMMC1	4	No	3.3, 1.8	208	104	SD/SDIO
SDMMC2	8	Yes	1.8	266, 208	533, 104	SDIO for secondary modem or eMMC device
SDMMC3	4	No	3.3, 1.8	208	104	SD/SDIO
SDMMC4	8	Yes	1.8	266	533	eMMC

3.3 Serial ATA (SATA) Controller

Standard	Notes
<i>Serial ATA Revision 3.1</i>	Including all errata, ENC, and TP, except DHU (direct head unload)
<i>Serial ATA Advanced Host Controller Interface (AHCI) Specification, Rev 1.3.1</i>	

The SATA controller enables a control path from the module to an external SATA device. A SSD / HDD / ODD drive can be connected. Controller can support the maximum throughput of a Gen 2 drive.

Features:

- Port multiplier support: command based switching (CBS)
- Supported Cables and connectors
 - Standard internal connector
 - Internal micro connector
 - Internal slimline connector
 - mSATA connector
 - BGA SSD interface
- Not supported: External connector (eSATA), USM, Internal LIF-SATA, DevSleep

3.4 Display Interfaces

The Tegra Display Controller Complex integrates two MIPI-DSI interfaces and two Serial Output Resources (SOR) to collect pixels from the output of the display pipeline, format/encode them to desired format, and then streams to various output devices. The SOR consists of several individual resources which can be used to interface with different display devices such as HDMI, DP or eDP.

3.4.1 MIPI Display Serial Interface (DSI)

The Display Serial Interface (DSI) is a serial bit-stream replacement for the parallel MIPI DPI and DBI display interface standards. DSI reduces package pin-count and I/O power consumption. DSI support enables both display controllers to connect to an external display(s) with a MIPI DSI receiver. The DSI transfers pixel data from the internal display controller to an external third-party LCD module.

Features:

- PHY Layer
 - Start / End of Transmission. Other out-of-band signaling
 - Per DSI interface: 1 Clock Lane; up to 4 Data Lanes
 - Supports link configuration – 1x4, 2x4
 - Supports dual link operation in 2x4 configurations for asymmetrical/symmetrical split in both left-right side or odd-even group split schemes.
 - DSC link compression
 - Maximum link rate 1.5Gbps as per MIPI D-PHY 1.1v version
 - Maximum 10MHz LP receive rate
- Lane Management Layer with Distributor
- Protocol Layer with Packet Constructor
- Supports MIPI DSI 1.0.1v version mandatory features
- Command Mode (One-shot) with Host and/or display controller as master
- Clocks
 - Bit Clock : Serial data stream bit-rate clock
 - Byte Clock : Lane Management Layer Byte-rate clock
 - Application Clock: Protocol Layer Byte-rate clock.
- Error Detection / Correction
 - ECC generation for packet Headers
 - Checksum generation for Long Packets
- Error recovery
- High Speed Transmit timer
- Low Power Receive timer
- Turnaround Acknowledge Timeout

3.4.2 High-Definition Multimedia Interface (HDMI) and DisplayPort (DP) Interfaces

Standard	Notes
<i>High-Definition Multimedia Interface (HDMI) Specification, version 2.0</i>	> 340MHz pixel clock Scrambling support Clock/4 support (1/40 bit-rate clock)
<i>DisplayPort 1.2a</i>	
<i>High-bandwidth Digital Content Protection (HDCP) System Specification, version 2.2</i>	
<i>High-bandwidth Digital Content Protection (HDCP) System Specification, version 1.3</i>	

The HDMI and DP interfaces share the same set of interface pins. A new transport mode was introduced in HDMI 2.0 to enable link clock frequencies greater than 340MHz and up to 600MHz. For transfer rates above 340MHz, there are two main requirements:

- All link data, including active pixel data, guard bands, data islands and control islands must be scrambled.
- The TMDS clock lane must toggle at CLK/4 instead of CLK. Below 340MHz, the clock lane toggles as normal (independent of the state of scrambling).

Features:

- On-chip HDCP key storage, no external Secure ROM required
- Support for both HDCP 1.3 and HDCP 2.2
- HDMI
 - HDMI 2.0 mode (3.4Gbps < data rate ≤ 6Gbps)
 - HDMI 1.4 mode (data rate ≤ 3.4Gbps)
 - Multi-channel audio from HDA controller, up to 8 channels 192kHz 24-bit.
 - Vendor Specific Info-frame (VSI) packet transmission
 - 24-bit RGB and 24-bit YUV444 (HDMI) pixel formats
 - Transition Minimized Differential Signaling (TMDS) functional up to 340MHz pixel clock rate
- DisplayPort
 - Display Port mode: interface is functional up to 540MHz pixel clock rate (i.e., 1.62GHz for RBR, 2.7GHz for HBR, and 5.4GHz for HBR2).
 - 8b/10b encoding support
 - External Dual Mode standard support
 - Audio streaming support

3.4.3 Embedded DisplayPort (eDP) Interface

Standard	Notes
<i>Embedded DisplayPort 1.4</i>	Supported eDP 1.4 features: <ul style="list-style-type: none"> ▪ Additional link rates ▪ Enhanced framing ▪ Power sequencing ▪ Reduced aux timing ▪ Reduced main voltage swing

eDP is a mixed-signal interface consisting of 4 differential serial output lanes and 1 PLL. This PLL is used to generate a high frequency bit-clock from an input pixel clock enabling the ability to handle 10-bit parallel data per lane at the pixel rate for the desired mode. Embedded DisplayPort (eDP) mode: 1.6GHz for RBR, 2.16GHz, 2.43GHz, 2.7GHz for HBR, 3.42GHz, 4.32GHz and 5.4GHz for HBR2.

NOTE: eDP has been tested according to DP1.2b PHY CTS even though eDPv1.4 supports lower swing voltages and additional intermediate bit rates. This means the following nominal voltage levels (400mV, 600mV, 800mV, 1200mV) and data rates (RBR, HBR, HBR2) are tested. This interface can be tuned to drive lower voltage swings below 400mV and can be programmed to other intermediate bit rates as per the requirements of the panel and the system designer.

The eDP block collects pixels from the output of the display pipeline, formats/encodes them to the eDP format, and then streams them to various output devices. It drives local panels only (does not support an external DP port), includes a small test pattern generator and CRC generator.

3.5 Audio Interfaces: I2S, PCM, TDM

Tegra X1 series processors support up-to four I2S interfaces. The I2S Controller transports streaming audio data between system memory and an audio codec. The controller supports I²S format, Left-justified Mode format, Right-justified Mode format, and DSP mode format, as defined in the Philips inter-IC-sound (I²S) bus specification. The timing in the following sections applies to any of these interfaces depending on whether they are configured for I2S/PCM or TDM mode.

The I2S and PCM (master and slave modes) interfaces support clock rates up to 24.5760MHz.

The I2S controller supports point-to-point serial interfaces for the I²S digital audio streams. I²S-compatible products, such as compact disc players, digital audio tape devices, digital sound processors, and those with digital TV sound may be directly connected to the I²S controller. The controller also supports the PCM and telephony mode of data-transfer. Pulse-Code-

Modulation (PCM) is a standard method used to digitize audio (particularly voice) patterns for transmission over digital communication channels. The Telephony mode is used to transmit and receive data to and from an external mono CODEC in a slot-based scheme of time-division multiplexing. The I2S controller supports bidirectional audio streams and can operate in half-duplex or full-duplex mode.

Features:

- Basic I2S modes to be supported (I2S, RJM, LJM and DSP) in both Master and Slave modes.
- PCM mode with short (one-bit-clock wide) and long-fsync (two bit-clocks wide) in both master and slave modes.
- NW-mode with independent slot-selection for both Tx and Rx
- TDM mode with flexibility in number of slots and slot(s) selection.
- Capability to drive-out a High-z outside the prescribed slot for transmission
- Flow control for the external input/output stream.

3.6 USB Interfaces

Standard	Notes
<i>Universal Serial Bus Specification Revision 3.0</i>	Refer to specification for related interface timing details.
<i>Universal Serial Bus Specification Revision 2.0</i>	USB Battery Charging Specification, version 1.0; including Data Contact Detect protocol Modes: Host and Device Speeds: Low, Full, and High Refer to specification for related interface timing details.
<i>Enhanced Host Controller Interface Specification for Universal Serial Bus revision 1.0</i>	Refer to specification for related interface timing details.
<i>eXtensible Host Controller Interface for Universal Serial Bus revision 1.0</i>	xHCI. Refer to specification for related interface timing details.

3.7 PCI Express (PCIe) Interface

Standard	Notes
<i>PCI Express Base Specification Revision 3.0</i>	Tegra processors meet the timing requirements for the Gen2 (5.0 GT/s) data rates. Refer to specification for complete interface timing details.

The Tegra X1 SoC supports two root port controllers each with x4 and x1 maximum lane width support

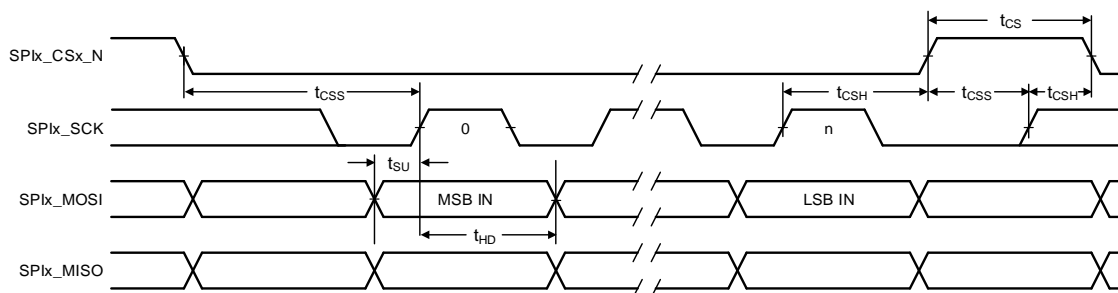
3.8 Serial Peripheral Interface

The SPI controllers operate up to 65 Mbps in master mode and 45 Mbps in slave mode. It allows a duplex, synchronous, serial communication between the controller and external peripheral devices. It consists of 4 signals, SS_N (Chip select), SCK (clock), MOSI (Master data out and Slave data in) and MISO (Slave data out and master data in). The data is transferred on MOSI or MISO based on the data transfer direction on every SCK edge. The receiver always receives the data on the other edge of SCK.

Features:

- Independent RX FIFO and TX FIFO.
- Software controlled bit-length supports packet sizes of 1 to 32 bits.
- Packed mode support for bit-length of 7 (8-bit packet size) and 15 (16-bit packet size).
- SS_N can be selected to be controlled by software, or it can be generated automatically by the hardware on packet boundaries.
- Receive compare mode (controller listens for a specified pattern on the incoming data before receiving the data in the FIFO).

- Simultaneous receive and transmit supported
- Supports Master and Slave modes of operation

Figure 2 SPI Master Timing Diagram

Table 5 SPI Master Timing Parameters

Symbol	Parameter	Min	Typ	Max	Unit
Fsck	SPIx_SCK clock frequency			65	MHz
Psck	SPIx_SCK period	1/Fsck			ns
t _{CH}	SPIx_SCK high time	50%Psck -10%		50%Psck +10%	ns
t _{CL}	SPIx_SCK low time	50%Psck -10%		50%Psck +10%	ns
t _{CRT}	SPIx_SCK rise time (slew rate)	0.1			V/ns
t _{CFT}	SPIx_SCK fall time (slew rate)	0.1			V/ns
t _{SU}	SPIx_MOSI setup to SPIx_SCK rising edge	2			ns
t _{HD}	SPIx_MOSI hold from SPIx_SCK rising edge	3			ns
t _{CSS}	SPIx_CSx_N setup time	2			ns
t _{CSH}	SPIx_CSx_N hold time	3			ns
t _{CS}	SPIx_CSx_N high time	10			ns

Note: Polarity of SCLK is programmable. Data can be driven or input relative to either the rising edge (shown above) or falling edge.

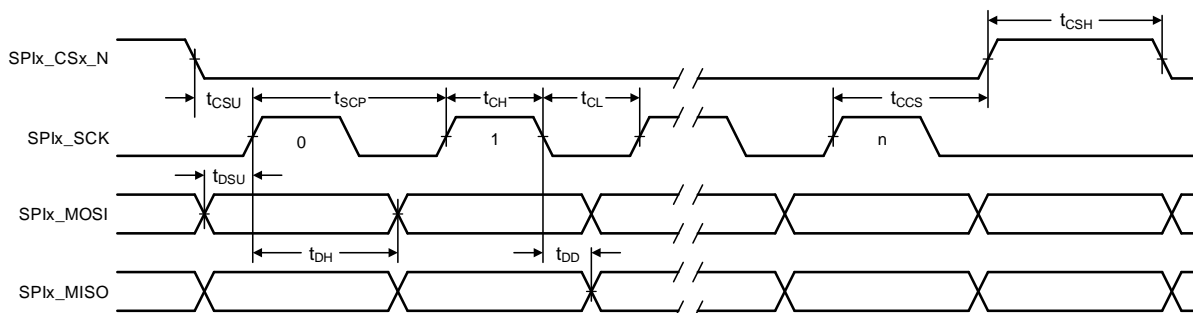
Figure 3 SPI Slave Timing Diagram


Table 6 SPI Slave Timing Parameters

Symbol	Parameter	Min	Typ	Max	Unit
t_{SCP}	SPIx_SCK period	$2 \cdot (t_{SDD} + t_{MSU}^1)$			ns
t_{SCH}	SPIx_SCK high time	$t_{SDD} + t_{MSU}^1$			ns
t_{SCL}	SPIx_SCK low time	$t_{SDD} + t_{MSU}^1$			ns
t_{SCSU}	SPIx_CSx_n setup time	1			t_{SCP}
t_{SCSH}	SPIx_CSx_n high time	1			t_{SCP}
t_{SCCS}	SPIx_SCK rising edge to SPIx_CSx_n rising edge	1		1	t_{SCP}
t_{SDSU}	SPIx_MOSI setup to SPIx_SCK rising edge	1		1	ns
t_{SDH}	SPIx_MOSI hold from SPIx_SCK rising edge	2		11	ns
t_{SDD}	SPIx_MISO delay from SPIx_SCLK falling edge (ALT1 ²)	3.5		16	ns
t_{SDD}	SPIx_MISO delay from SPIx_SCLK falling edge (ALT2 ²)	3		13	ns
t_{SDD}	SPIx_MISO delay from SPIx_SCLK falling edge (ALT3 ²)	4		17	ns

1. t_{MSU} is the setup time required by the external master

2. ALT1/2/3 refers to the position of the SPI pins in the Signal Pinout Multiplexing tables in Section 3.1, *Signal List and Multiplexing Functions*.

Note: Polarity of SCLK is programmable. Data can be driven or input relative to either the rising edge (shown above) or falling edge.

3.9 Inter-Chip Communication (I2C) Controller

Standard	Notes
NXP inter-IC-bus (I ² C) specification	

The I2C controller implements an I²C-bus specification compliant I2C master and a slave controller. The I2C controller supports multiple masters and slaves in: Standard-mode (up to 100Kbit/s), Fast-mode (up to 400 Kbit/s), Fast-mode plus (Fm+, up to 1Mbit/s) and High-speed mode (up to 3.4Mbit/s) of operations. A general purpose I2C controller allows system expansion for I2C-based devices, such as AM/FM radio, remote LCD display, serial ADC/DAC, and serial EPROMs, as defined in the NXP inter-IC-bus (I²C) specification. The I2C bus supports serial device communications to multiple devices. The I2C controller handles bus mastership with arbitration, clock source negotiation, speed negotiation for standard and fast devices, and 7-bit and 10-bit slave address support according to the I2C protocol and supports master and slave mode of operation.

3.10 UART Controller

UART controller provides serial data synchronization and data conversion (parallel-to-serial and serial-to-parallel) for both receiver and transmitter sections. Synchronization for serial data stream is accomplished by adding start and stop bits to the transmit data to form a data character. Data integrity is accomplished by attaching a parity bit to the data character. The parity bit can be checked by the receiver for any transmission bit errors.

Features:

- Synchronization for the serial data stream with start and stop bits to transmit data and form a data character
- Supports both 16450- and 16550-compatible modes. Default mode is 16450
- Device clock up to 200MHz, baud rate of 12.5Mbits/second
- Data integrity by attaching parity bit to the data character
- Support for word lengths from five to eight bits, an optional parity bit and one or two stop bits
- Support for modem control signals

- DMA capability for both TX and RX
- 8-bit x 36 deep TX FIFO
- 11-bit x 36 deep RX FIFO. 3 bits of 11 bits per entry will log the RX errors in FIFO mode (break, framing and parity errors as bits 10,9,8 of FIFO entry)
- Auto sense baud detection
- Timeout interrupts to indicate if the incoming stream stopped
- Priority interrupts mechanism
- Flow control support on RTS and CTS
- Internal loopback
- SIR encoding/decoding (3/16 or 4/16 baud pulse widths to transmit bit zero)

3.11 Video Input Interfaces

3.11.1 MIPI Camera Serial Interface (CSI)

Standard	Notes
MIPI CSI 2.0 Receiver specification	

The Tegra X1 SoC incorporates three MIPI CSI x4 bricks supporting a variety of device types and camera configurations. The Camera Serial Interface (CSI) is based on MIPI CSI 2.0 standard specification and implements the CSI receiver which receives data from an external camera module with a CSI transmitter.

Features:

- Supported camera configurations:
 - 1 x4: single camera with a 4 lane sensor using any one of three MIPI x4 bricks.
 - 2 x4: stereo pair with 4 lanes for each camera using any pair of three MIPI x4 bricks.
 - 2 x2: dual camera mode, breaking up any MIPI x4 brick to two x2 sub blocks; can support up to 6 camera streams simultaneously.
- Supported input data formats:
 - RGB: RGB888, RGB666, RGB565, RGB555, RGB444
 - YUV: YUV422-8b, YUV420-8b (legacy), YUV420-8b, YUV444-8b
 - RAW: RAW6, RAW7, RAW8, RAW10, RAW12
 - DPCM: user defined
 - User defined: JPEG8
 - Embedded: Embedded control information
- Supports single-shot mode
- D-PHY Modes of Operation
 - High Speed Mode – High speed differential signaling up to 1.5Gbps; burst transmission for low power
 - Low Power Control – Single-ended 1.2V CMOS level. Low speed signaling for handshaking.
 - Low Power Escape –Low speed signaling for data, used for escape command entry only. 20Mbps

If the two streams come from a single source, then the streams are separated using a filter indexed on different virtual channel numbers or data types. In case of separation using data types, the normal data type is separated from the embedded data type. Since there are only two pixel parsers, virtual channel and embedded data capability cannot be used at the same time.



3.11.2 Camera / VI (Video Input)

The Video Input (VI) is the cross bar unit for camera processing. It receives video from CSI and directs that data to either main memory or ISP.

3.12 Debug Interface

The Tegra X1 SoC has an optional JTAG interface that can be used for SCAN testing or for communicating with integrated CPU.

Notice

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation ("NVIDIA") does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regard to the purchase of the NVIDIA product referenced in this specification.

NVIDIA products are not designed, authorized or warranted to be suitable for use in medical, military, aircraft, space or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on these specifications will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this specification. NVIDIA does not accept any liability related to any default, damage, costs or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this specification, or (ii) customer product designs.

No license, either express or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this specification. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA 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 rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA. Reproduction of information in this specification is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

ARM

ARM, AMBA and ARM Powered are registered trademarks of ARM Limited. Cortex, MPCore and Mali are trademarks of ARM Limited. All other brands or product names are the property of their respective holders. "ARM" is used to represent ARM Holdings plc; its operating company ARM Limited; and the regional subsidiaries ARM Inc.; ARM KK; ARM Korea Limited.; ARM Taiwan Limited; ARM France SAS; ARM Consulting (Shanghai) Co. Ltd.; ARM Germany GmbH; ARM Embedded Technologies Pvt. Ltd.; ARM Norway, AS and ARM Sweden AB

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA, the NVIDIA logo, and Tegra are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2014 – 2016 NVIDIA Corporation. All rights reserved