

Optimal Systems Architecture for High Definition Video

Ajit V Rao
ajitr@ti.com
Texas Instruments

Agenda

- Video in the Infrastructure: The triggers
- Trends in Key areas – Surveillance, Conferencing, Streaming
- Key system care-about
- Power & Cost / Channel Trends
- Video System-on-Chip architecture trends
- The confluence & road ahead
- Conclusions

Video in the Infrastructure: The Triggers

- Massive expansion in content generated and available
 - User & professional
 - Entertainment, Education, Live content, Surveillance
 - Network storage
- Large number of video capable end-points
 - Powered by low-cost, low-power System-on-Chips (SoC)
 - “Three screens”: PCs, TVs and handsets
 - IP Connectivity rapidly extending to TVs and handsets
 - Large variance in bandwidth, power constraints, protocols / formats / codecs / resolutions supported
- Continuous growth in resolution/fps – HD everywhere, moving to 3D
- Record locally, Store Locally → Networked, any-time accessible, optimized search/ delivery
- Broadcast → IPTV → Streaming → Surveillance, Conferencing

Broadcast / Cable / IPTV Infrastructure

- Broadcast / IPTV Encoders:
 - Multi-channel, Multi-format HD / SD
 - MPEG2 moving to H.264 and SVC
 - Lower form factors
 - Multi-format outputs
- Broadcast Decoders
- Video Quality Analyzers

Video Surveillance - Trends

- From: Analog end-points recorded on Digital Video Recorders (HDD storage) to: IP Network cameras (Network storage)
- Custom Cabling → IP Networked w. PoE/ WiFi
 - Impacts network design, storage, power design of systems
- Hybrid (Analog & Digital) DVRs are key in the intermediate.
- Standard Definition ~ 1Mbps → HD (1080p and higher) @ 10 Mbps; Digital Zoom.
- Codecs: H.264, SVC
- Data storage & analysis moves to the network, redundant storage (RAID).
- Unified management of archiving
- End-points include basic analytics: “Edge analytics”. Analytics servers on the network.
- Combined into intelligent / actionable information on the network - alarms, event management.
- Stronger data security capabilities – encrypted data for storage
- Interoperability standards across vendors (ONVIF/PSIA) for reducing system costs.
- Multi-channel Viewing stations with advanced features - video texturing, camera / event flows, low latency

Video Conferencing - Trends

- Higher definitions and multi-screen systems (Tele-presence)
- VC systems moving from “closed point-point” to Open standards-compliant systems: H.323 / SIP / XMPP
- Codecs – H.264, SVC: 128 kbps → 12Mbps.
- Inter-operability across vendors and end-points: Tele-presence, Desktop, WebEx, cell-phones
- End-points talk to “Multipoint Conferencing Units” in the network
- MCUs trans-code to multiple standards / resolutions - handling end-points with vastly different capabilities.
- Low latency is a major care-about
- MCUs require multichannel decode, scale, composition, graphics overlay, re-encode - each end-point may require a separate composition of the conference participants
- Scalable video coding has several advantages

Video Streaming - Trends

- Long tail content growing at incredible rate
 - Content indexing and search critical
- Live/TV content, Video sharing very popular – tighter integration with social networking
- Content is moving to HD / 3D.
- Lot of content generated on the handset.
- Content consumed on multiple devices with vastly different capabilities
- Monetization becoming key concern - Ad insertion / overlay
- 3G adoption is growing – lots of content will be consumed over unreliable wireless channels. Rate shaping for wireless channels needed for best user experience

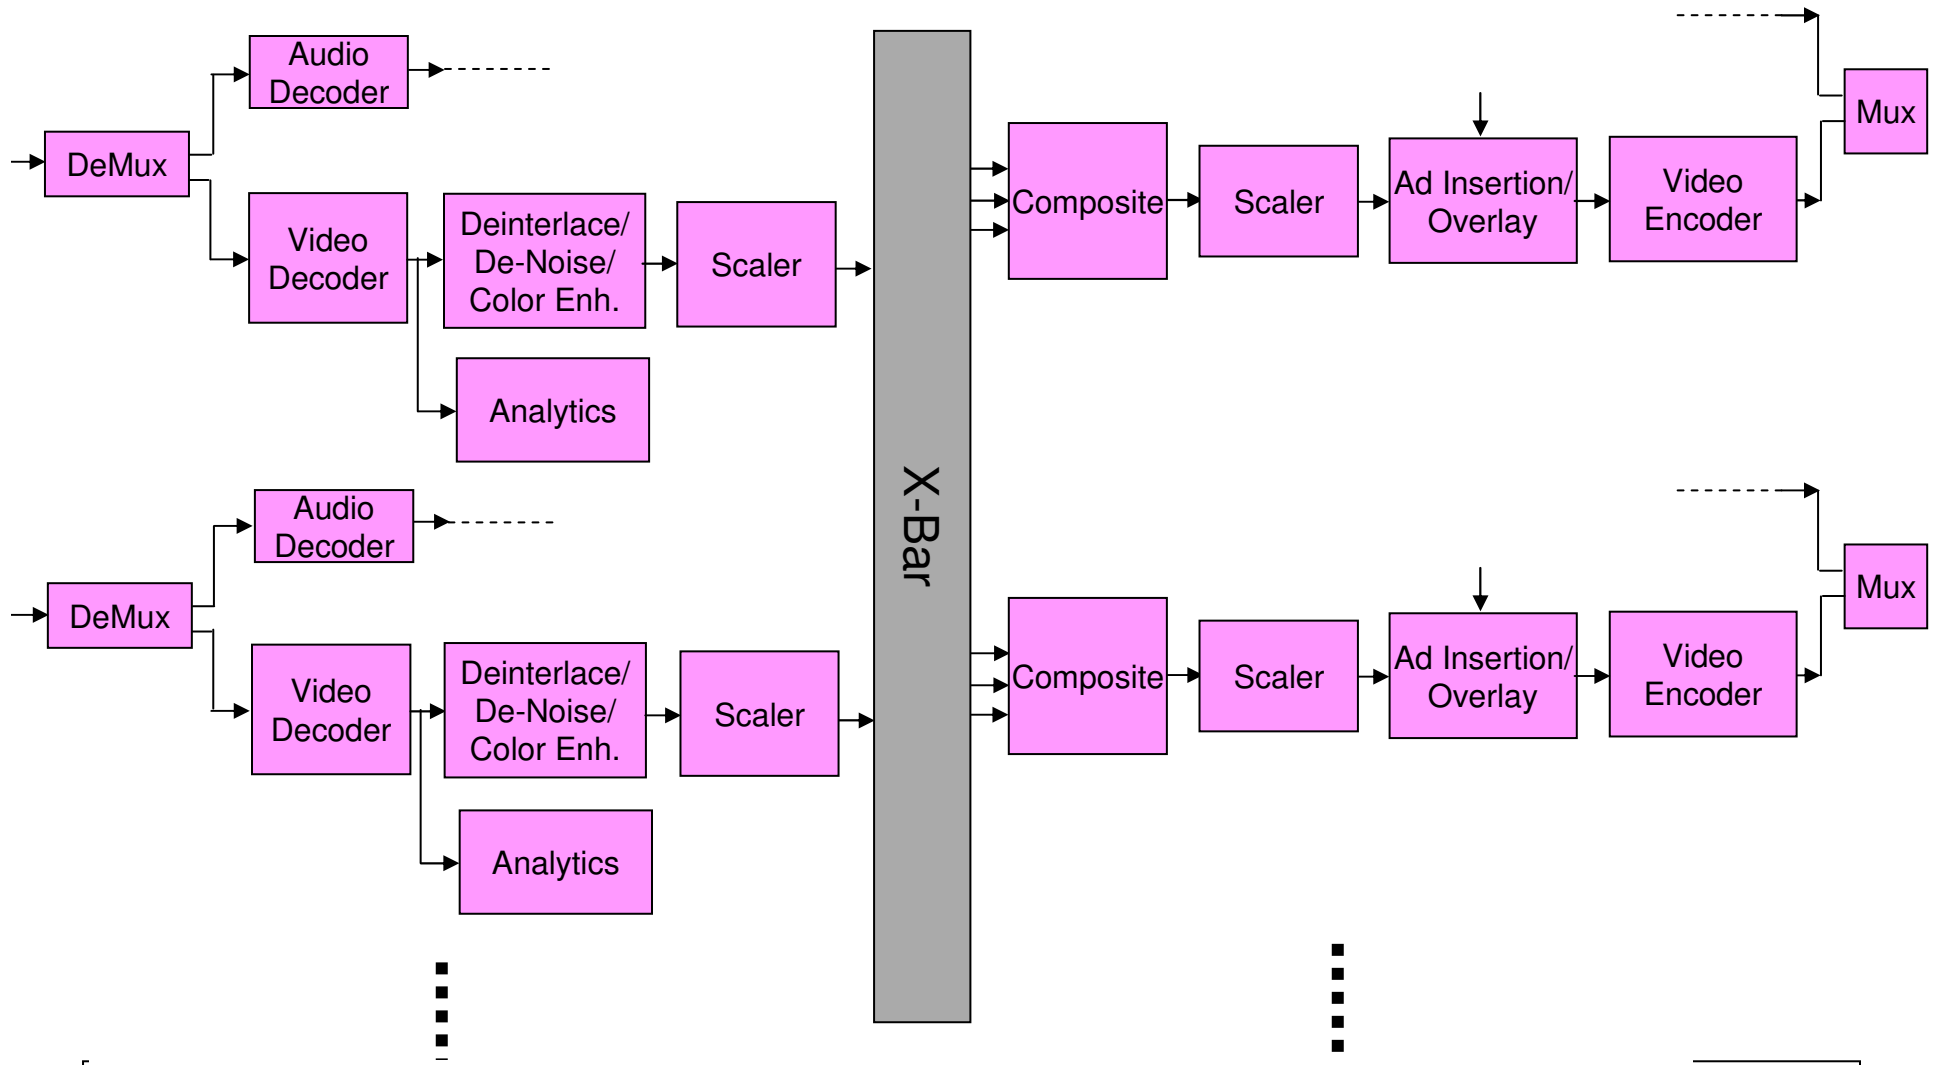
Challenges with Video Infrastructure

- Scalability
 - Expectations of 100s → 1000s of channels / system
 - Must support scalable number of lower resolution channels (low per-frame overheads)
 - Scalability with audio
- Reliability
 - Redundancy needed (often 1:1)
- Low power consumption
 - Power constraints will limit system capabilities
 - 2RU systems must be limited to < 200W
 - Similar power constraints for PCIe cards / mezzanine cards
 - Drives lower power/channel
- Low latency
 - Especially in 2-way calls
- Low cost / channel

Problem Statement

- High-density video infrastructure must handle several simultaneous channels of:
 - Trans-coding (protocol & format changes, resolution changes, bit-rate changes, rate shaping)
 - Video composition and encoding
 - Video analytics
 - Ad Insertion
- At the lowest power, latency and cost / channel and down-time

Typical Video Processing Pipeline



(deliberately doesn't show non-video functions – storage, retrieval, networking etc.)

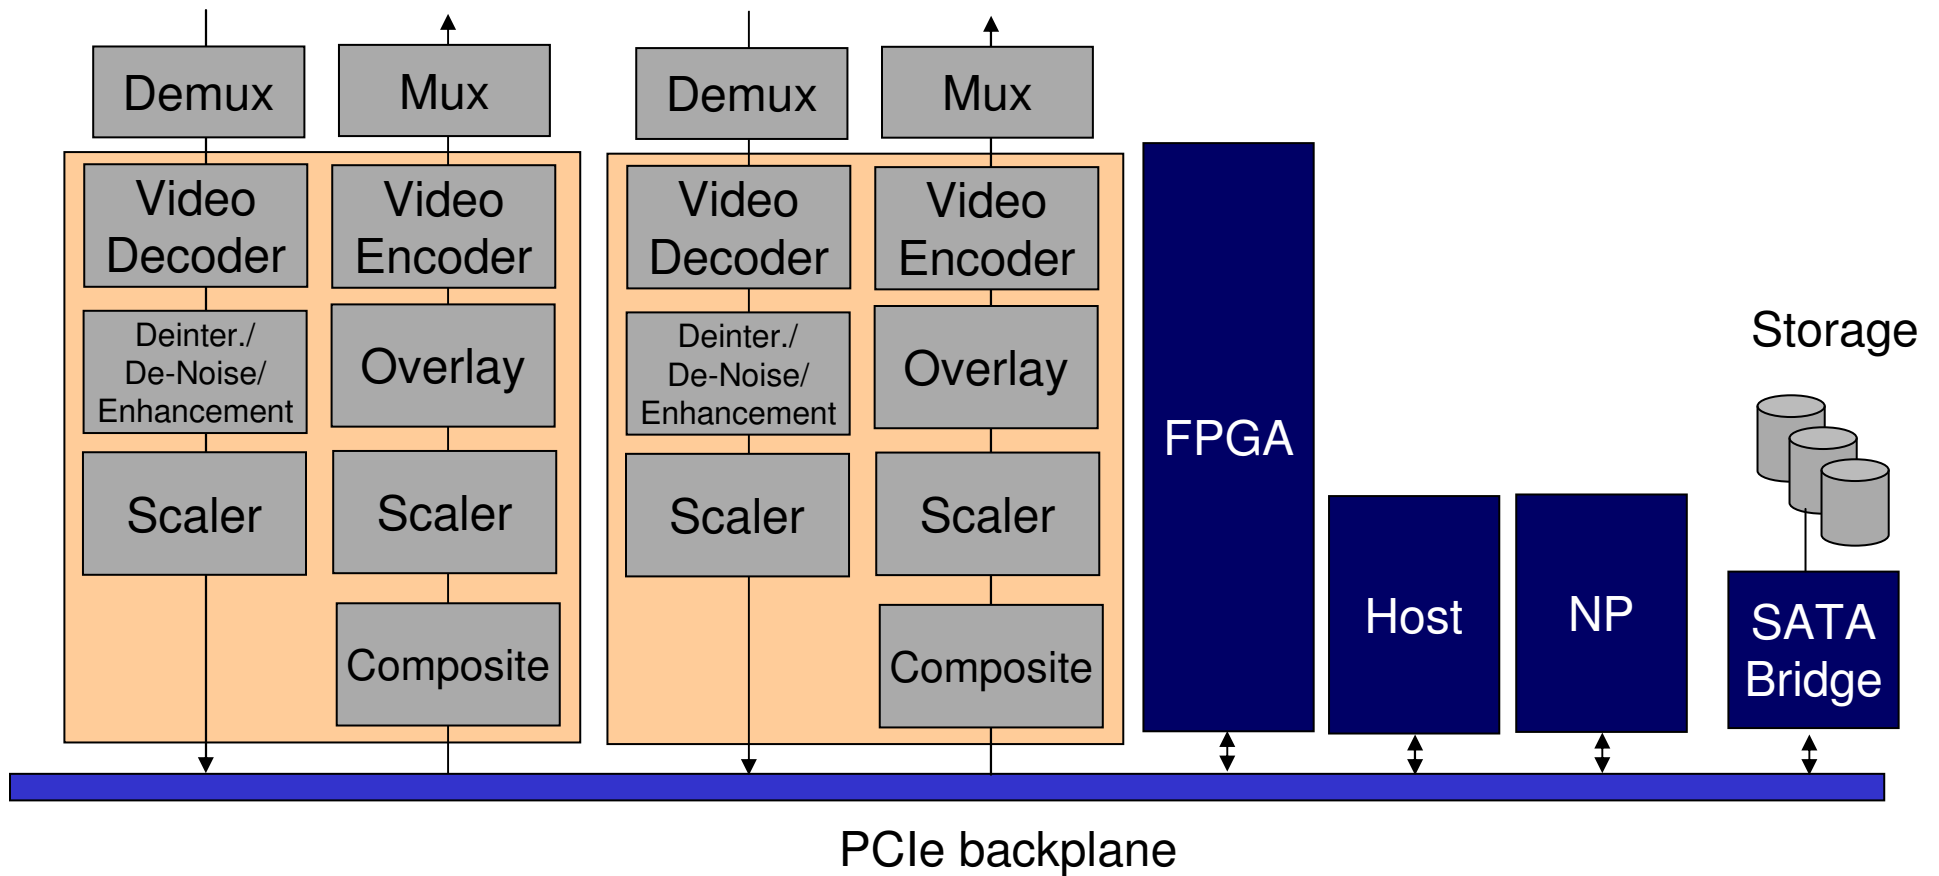
x86 Performance

- Standard off-the-shelf servers will scale poorly for video infrastructure applications.
 - Power consumption, cost per channel.
- Even with process-node scaling, cost and power is prohibitive
- Architectures optimized for video essential

Optimizing for video

- Optimal: Several SoCs sharing a common backplane (PCI / PCIe).
- SoC:
 - Optimized multi-format video compression engines
 - Programmable RISC / DSP processor
 - Optimized HW for resizing, de-noising and de-interlacing.
 - HW blending of video and graphics for overlays
 - Host processor w/ networking and storage
 - Video capture and display support (maybe)
- Optionally a separate host processor, a network processor and storage on backplane
- Custom or dedicated signal processing / other functions may be implemented on FPGAs that share the backplane.

Optimal System Architecture for high density video



Ideal Video SoC

- Intelligent partitioning across optimized processors / modules for different functions.
 - Common memory system ideal
- Memory throughputs requirements are high
 - Smart multi-level caching with dedicated memories
- Processors optimized for key video kernels
 - Motion and edge adaptive signal processing –
 - Scaling, De-interlacing, De-Noising
 - Requires algorithm-optimized custom-designed modules for highest efficiency
 - Video compression kernels must comply to multiple standards
 - Variable Length coding, Prediction & Interpolation Filters, Transforms and Loop Filters particularly challenge designs
 - Semi-programmable approach necessary for practical implementations
- One processor fits-all approach fails. Need solutions with:
 - Heterogeneous Multi-core
 - Function-optimized cores
 - Programmable pipelines with support for concurrency
 - Low frame level overheads
 - Common memory subsystem

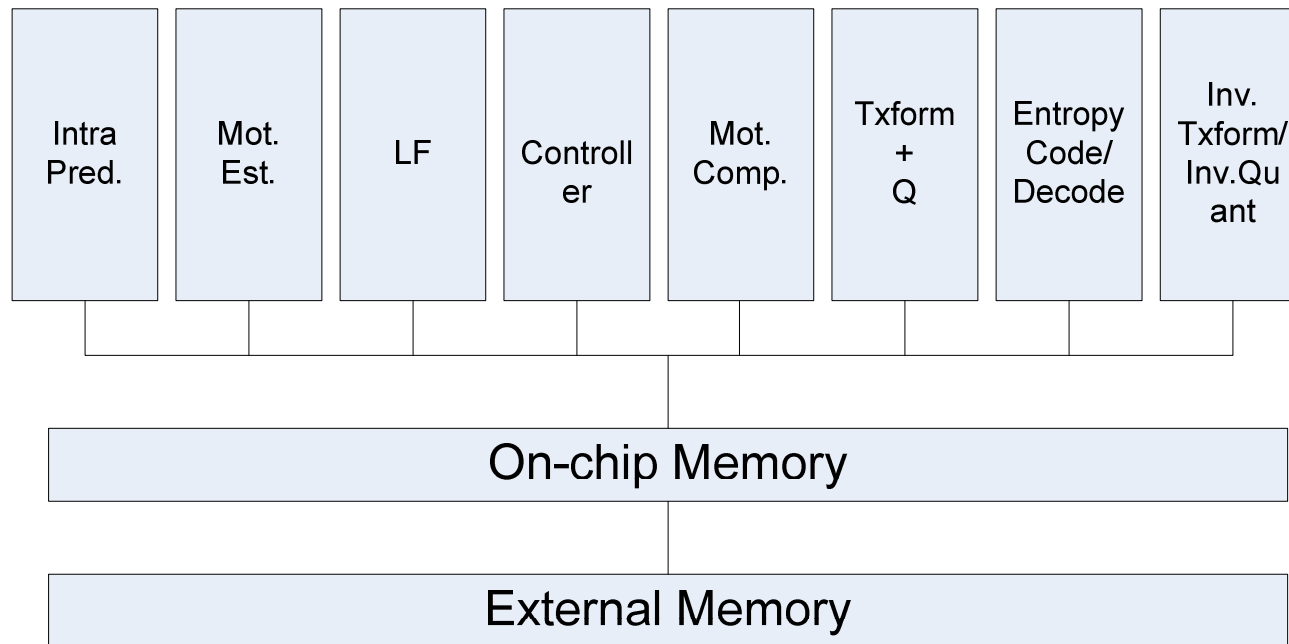
Options for compression engine

- Fully Programmable:
 - Flexible
 - Efficiency depends on processor architecture
 - Parallel compute
 - VLIW (ex: TI c6000)
 - SIMD (ex: Intel MMX, ARM Neon)
 - Multi-core (ex: Tiler) architectures.
 - Works for SD. Doesn't scale well to HD
 - High power consumption requirements (~10MIPS/mW). Typically ~ 300-500MHz / SD channel)
- Fully hardwired:
 - Work well for mature video codec standards
 - Large production volumes.
 - High cost of development and validation
 - High design cycle-time
 - Inflexible to changes in codec standards / system requirements.
 - Typical Power is 1000 MIPS/mW
- Hybrid architectures
 - Combination of hardwired / programmable

Hybrid Architectures - Motivation

- Well-designed combination of programmable processors and hard-wired IP blocks
- High performance and low power
- Multi-codec support
- Tunability to end-product needs.
- Function-specific HW cores or “co-processors”
 - Concurrently running in a pipeline
- Programmable processors/controllers
 - Programmability
 - Orchestrating co-processor pipeline.
- Examples: TI DaVinci and ST Nomadik solutions.

Hybrid Architecture for video compression engine



- Can deliver ~ 1000 MIPS/mW while offering tremendous implementation flexibility.

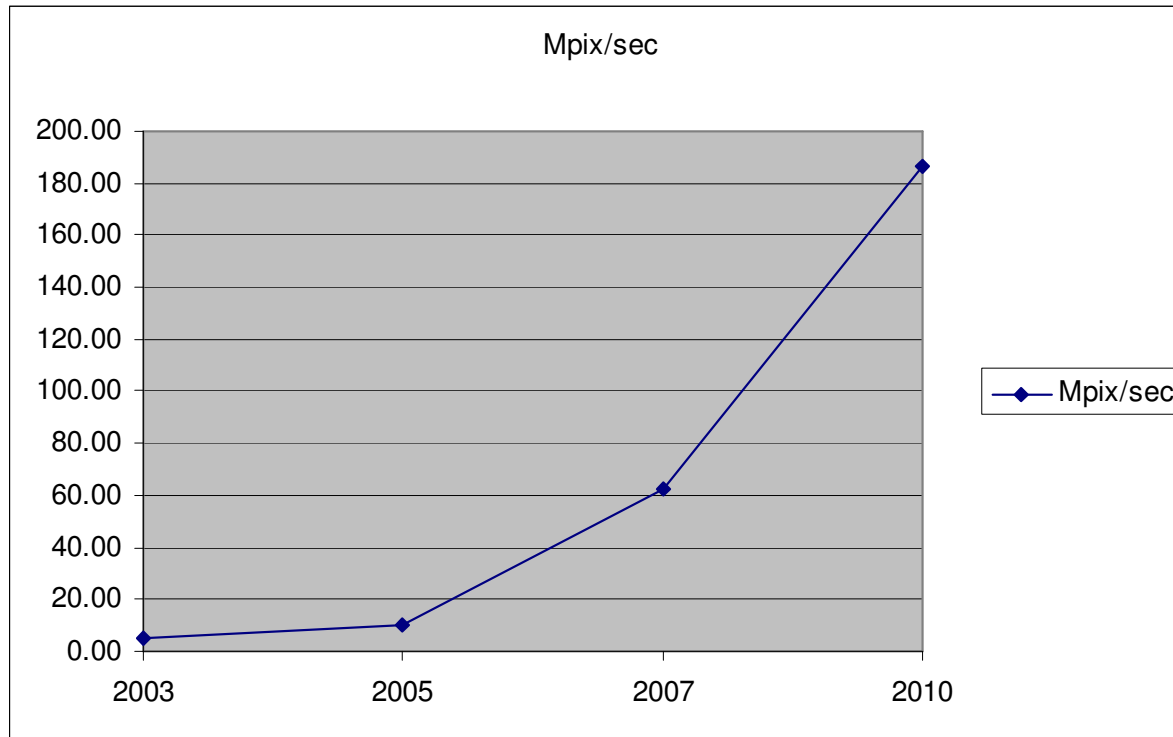
Video Processing functions - considerations

- Video Scaling:
 - Typically hardware optimized.
 - Edge adaptive scaling crucial for quality
 - Support for non-linear scaling when aspect ratio changes.
 - Programmable filter coefficients.
- De-interlacing
 - Spatio-temporal
 - Motion-adaptive / Motion-compensated
 - Support for film mode
 - Edge adaptation for spatial interpolation.

Memory architecture design

- Optimal memory architecture critical for optimized video SoCs.
- Memory bandwidth can be typically 20-100 bytes/pixel processed.
- Memory architecture must be well-balanced
 - Recommend two-layer architecture with large internal memory and efficient functional pipeline to reduce off-chip fetching.
- Input and reference pictures stored off-chip
- Transferred JIT into small/fast local on-chip memory.
- All HW engines operate internal memory.
- Custom video-optimized DMA engines to handle on-chip/off-chip transfers.

Performance improvements in Video SoCs



	2003	2005	2007	2010
HW accel. Video codecs	Red	Green	Green	Green
HW accel. Video Scaling	Red	Red	Green	Green
Video-optimized memory systems	Red	Red	Yellow	Green
HW accel. De-noising, De-interlacing	Red	Red	Red	Green
HW assisted color enhancement	Red	Red	Red	Green
Gigabit Networking	Red	Red	Red	Green

Also power efficiency (mW/Mpix) and cost (\$/Mpix) continue to exponentially reduce

Conclusions

- High density video infrastructure continues to rapidly grow as a market
 - Driven by pervasive networked content and devices
- High density, low power, low latency, low cost major care-about
- Hybrid multi-core SoCs on a common backplane, with host and NP ideal solution compared to off-the-shelf.
- Costs and power at ideal points for take-off.