

The Intel® Quick Sync Video Technology in the 2nd-Generation Intel Core™ Processor Family

Dr. Hong Jiang
Sr. Principal Engineer
Chief Media Architect
Intel Corporation
August 19, 2011

Hot Chips 23, Palo Alto, CA

Legal Notices

This presentation is for informational purposes only. INTEL MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

For more complete information about performance and benchmark results, visit www.intel.com/benchmarks

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Atom, Centrino Atom Inside, Centrino Inside, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, IntelDX2, IntelDX4, IntelSX2, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, Viiv Inside, vPro Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

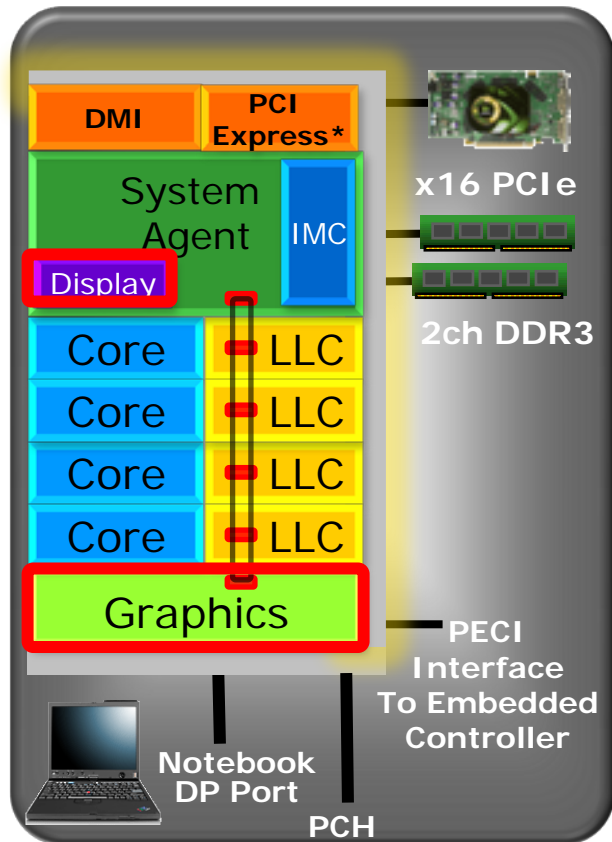
* Other names and brands may be claimed as the property of others.

Copyright © 2011, Intel Corporation. All rights reserved.

Agenda

- Processor Graphics Overview
- Intel® Quick Sync Video Technology Details
- Results and Conclusion

Processor Graphics in the 2nd Gen Intel Core™ Microprocessor



2nd Gen Intel Core™
Microprocessor
(codename Sandy Bridge)

- On-die Processor Graphics (PG)
 - Graphics Render Engines on the Ring
 - Display Engine in System Agent
- PG Integration Benefits
 - Graphics on the latest 32-nm process
 - Better frequency, power and area
 - Share high-BW Last Level Cache (LLC)
 - Up to 384 GB/s (4-Core), 192 GB/s (2-Core)*
 - Gfx driver controls per buffer LLC mapping
 - Dynamic Cores and PG power distribution
 - PG dynamic frequency up to 1350MHz
 - Fine-grain power management

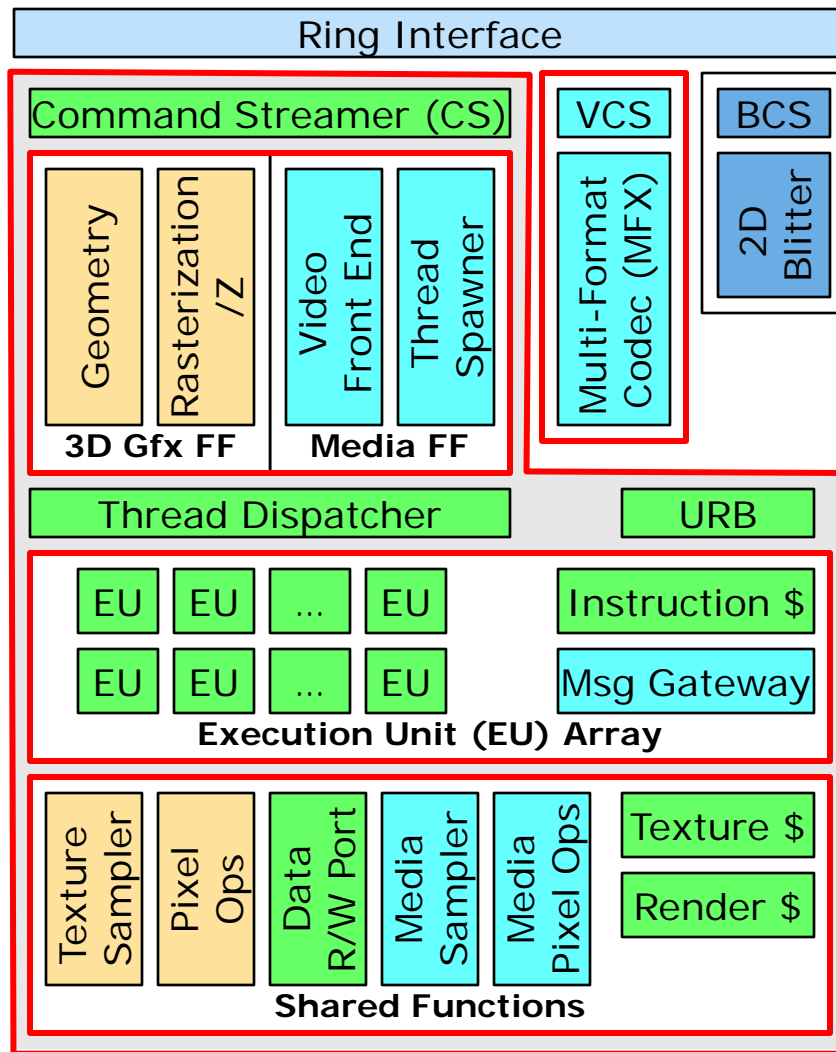
* At ring frequency of 3GHz

Processor Graphics: An Unified 3D Graphics and Media Architecture

The Main Render Engine: Unified for 3D Gfx and Media

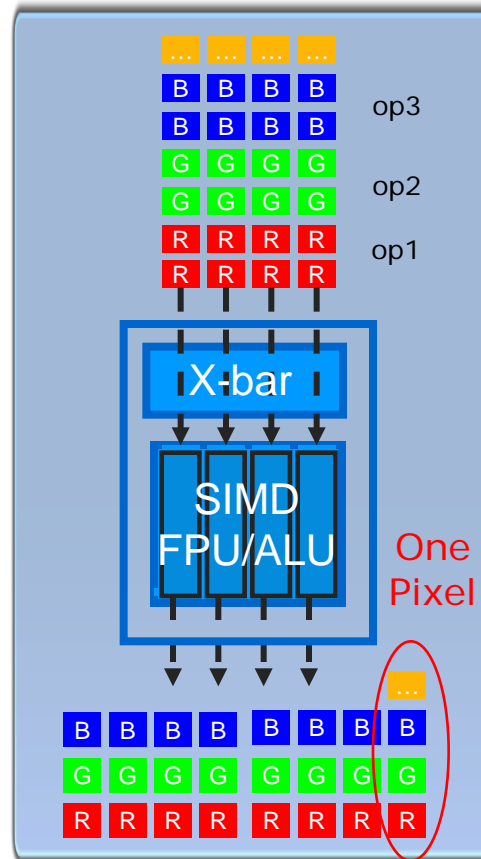
- Fixed Function (FF) Pipelines
 - Thread generation and control
 - 3D Gfx or Media FF controls EU array at a given time
- Execution Unit Array
 - Shared between 3D and Media
 - ISA optimized for both
- Shared Functions
 - Accelerators for filtered load, scatter & gather, filtered/blended store operations

MFX: A parallel codec engine runs on a separate context

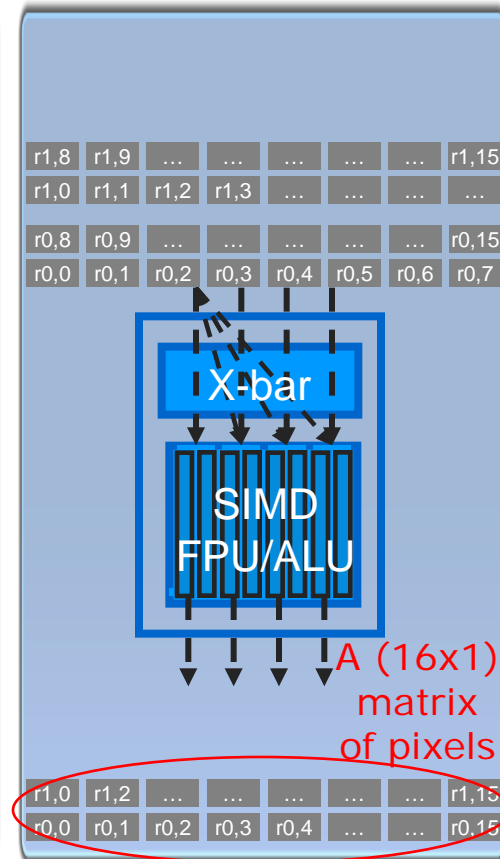


PG Execution Unit: Optimized for 3D Graphics and Media

- PG Execution Unit:
 - 5 HW threads per EU w/ zero-cost thread switching
 - 128 256-bit registers/thread
 - 128-bit FPU
- Optimization for 3D Graphics
 - FMAD, PLN, LRP, DP4/3/2
 - Per channel conditionals
 - MathBox per EU (sin, cos, exp)
- Optimization for Media
 - Byte/word/dword operations
 - Higher precision Accumulators
 - Full byte-xbar per src operand
 - Register 2D region addressing
 - Register indexed addressing



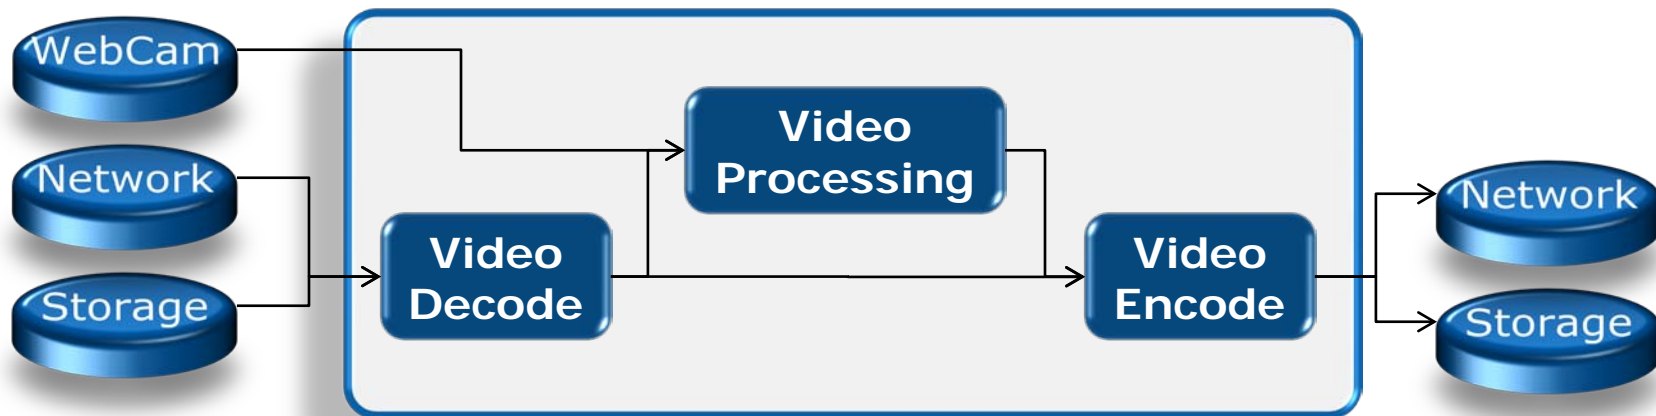
Channel Serial ISA
for 3D Pixel Shader



Vector/Matrix ISA
for Media

Intel® Quick Sync Video Technology

- Intel® Quick Sync Video Technology, in 2nd Gen Intel Core™ processor family, is hardware acceleration for fast media content creation, editing and conversion.
- It is enabled by a set of key media innovations for video decode, encode and processing in the Process Graphics



Intel® Quick Sync Video: Challenges

Challenges

- “Faster than Real-Time”
 - Significant boost over previous generations



- High visual quality
 - Quality cannot be compromised for the pursuit of speed



- Low power for on the go
 - Transcoding & editing no longer confine to desktop



- Low host CPU utilization for multi-tasking

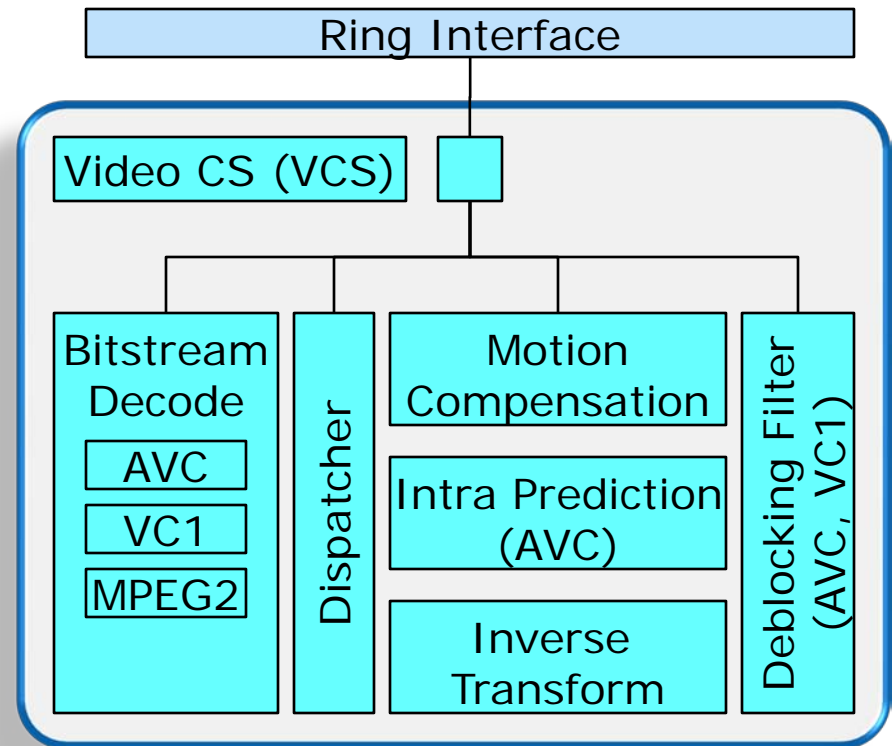


Approaches

- Performance boost through dynamic power delivery
- Multiple engines running in parallel to gain throughput
- Use hardware acceleration
- Programmable solution to enable algorithm innovations
- Significant use of fixed functions for Cdyn reduction
- “Race-to-halt” to minimize PG & overall processor power
- Full task offloading, which also reduces CPU/Gfx synchronization overhead

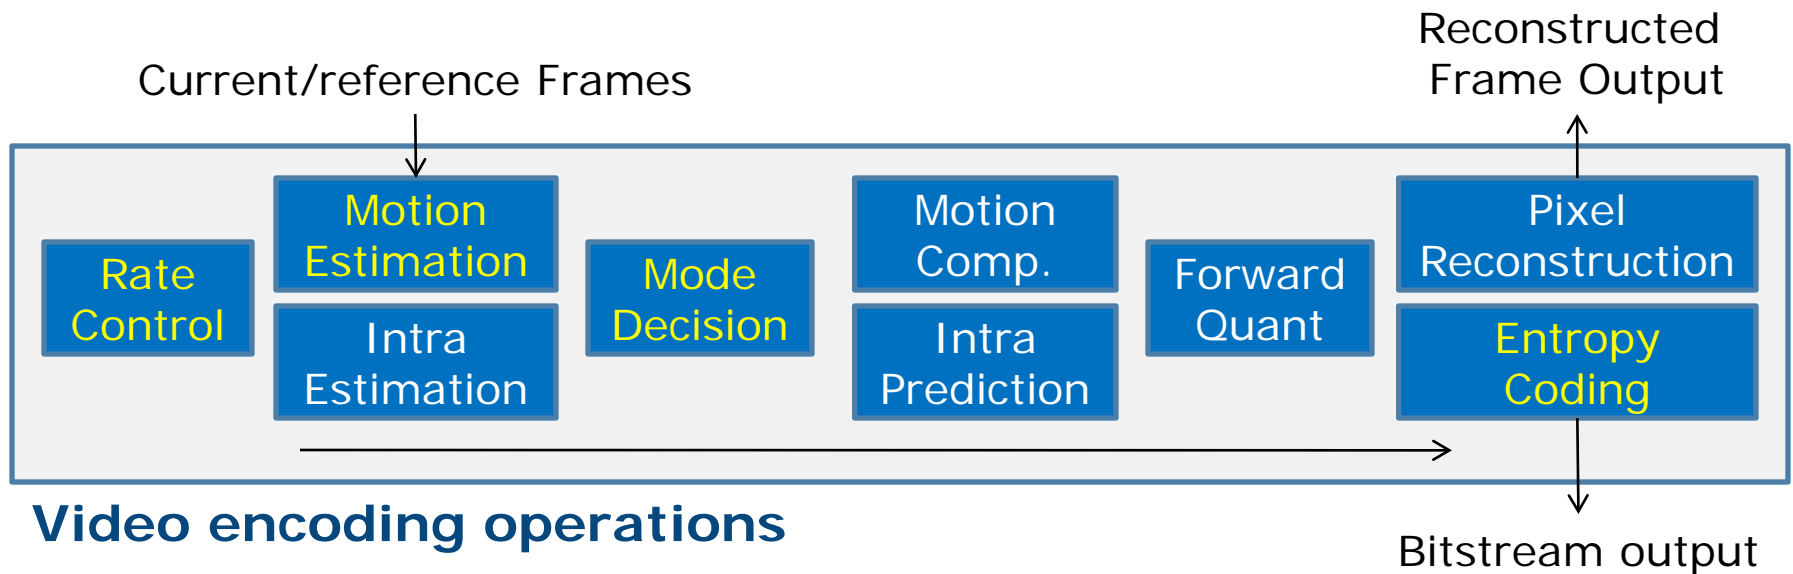
Video Decode: All Hardware Solution

- Multi-Format Codec (MFX) contains a full-hardware decoder
- Unified pipeline supporting multiple main stream video formats (BD and BD S3D)
 - AVC/MVC main & high profile
 - VC1 main & advanced profile
 - MPEG2 main profile
- Shared data path optimized for area & power
- State-less frame switch for multi-stream support
- Support PG's max dynamic frequency (up to 1350MHz)



MFX - Decoder Portion

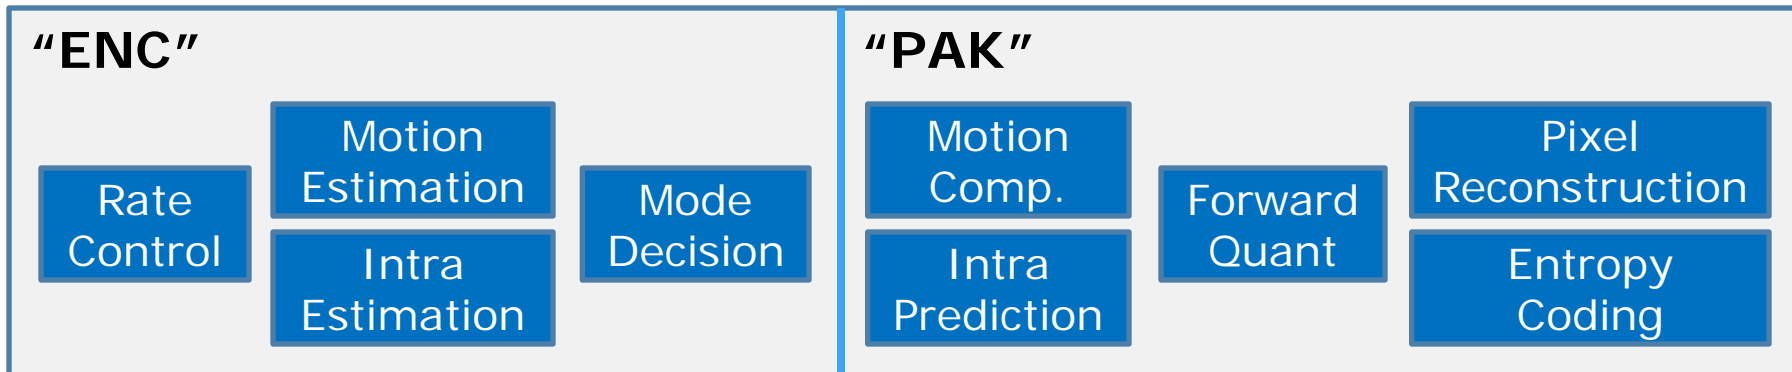
Understanding a Video Encoder



- Motion estimation is the most compute demanding stage
 - It also needs to be flexible and coupled with mode decision
- Programmability is critical for mode decision & rate control
- Entropy coding is a high performance serial compute
 - 30Mbps CABAC alone may require ~GOPS
 - Hard to map to pixel-oriented graphics processors

Video Encode: A Hybrid HW/SW Solution

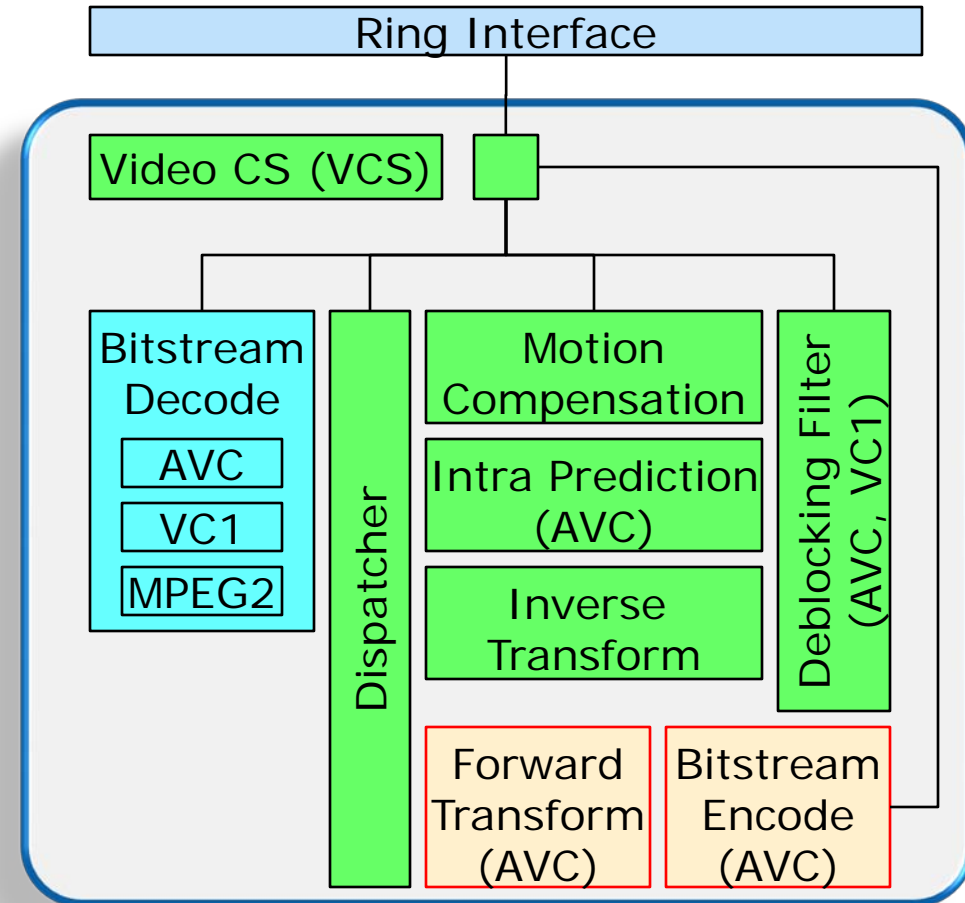
A 2-Stage Video Encoder:



- "ENC" maps to Programmable EU Array
 - Media Sampler does the heavy lifting for Motion Estimation
 - Flexible "ENC" enables algorithm tuning & feature additions
- "PAK" maps to MFX hardware pipeline
 - Reuse decoder's pixel reconstruction circuitry
- "ENC" and "PAK" may run concurrently on different frames for better throughput

Video Encode: "PAK" in MFX

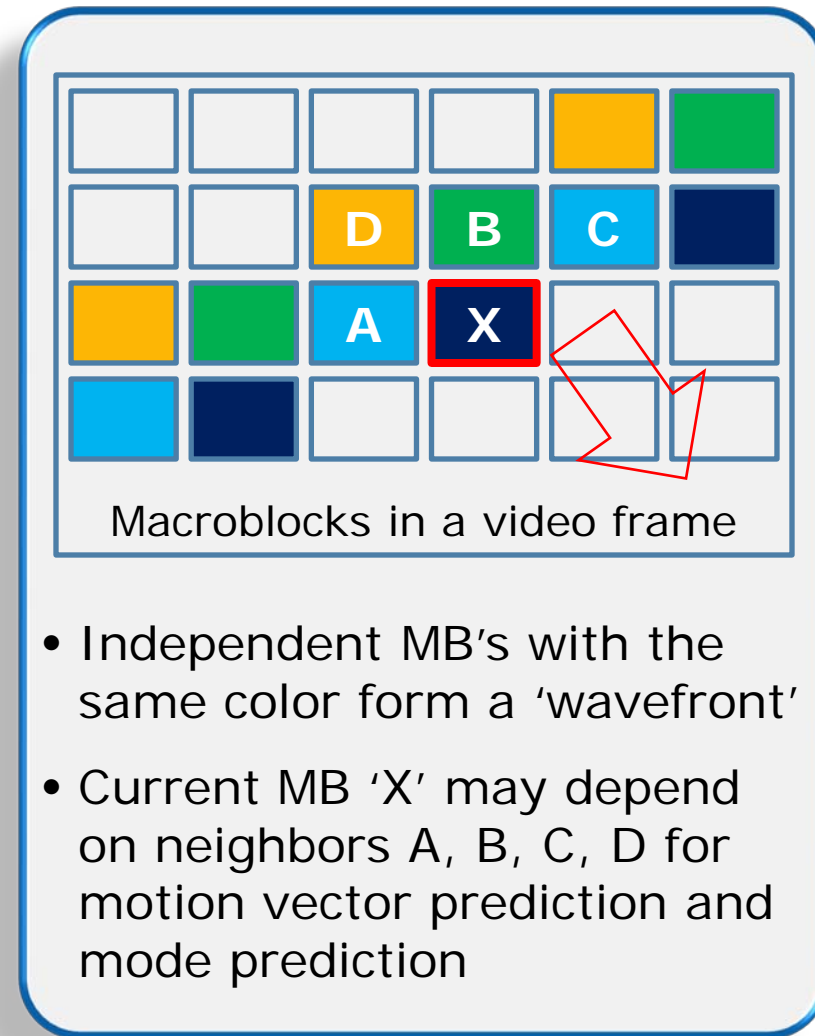
- MFX's configurable datapath supports encode PAK
- PAK hardware supports critical AVC features
 - CABAC and CAVLC
 - All major and minor partitions (16x16, 8x8 to 4x4)
 - Skip and Direct Modes
 - Transform 8x8 and 4x4
 - In-loop deblocking
 - AVC Spec & HRD conformance
 - ...
- Flexible control with MB-level PAK interface
- Shortened intra-prediction loop delivers high perf.



MFX – Full Codec Stack

Video Encode: "ENC" Parallel Programming

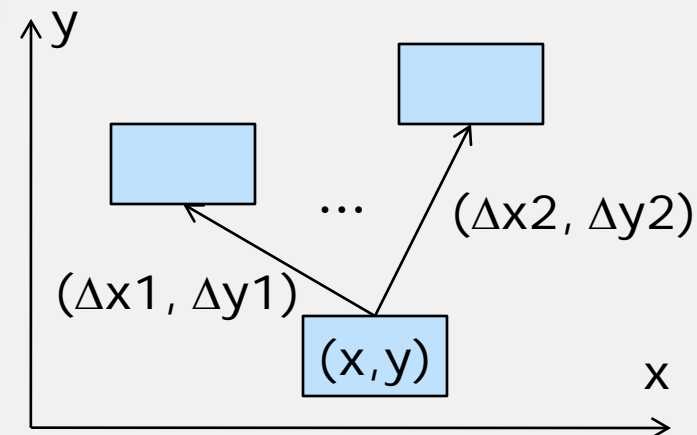
- Multi-threading to
 - utilize computes in EU Array, and
 - to cover Media Sampler DRAM latency
- 'Explicit' parallel programming, critical for an encoder:
 - One thread for one macroblock (MB)
 - A thread mixes pixel computation (rich vector/matrix data types), motion prediction (scalar) and mode decision (scalar)
 - Threads are dispatched in 'wavefronts'
 - Hardware scoreboard manages thread dependency/ordering
 - Shared memory and scoreboard apply to all threads in EU array



Video Encode: "ENC" Parallel Programming

- Hardware Scoreboard

- HW Scoreboard maintains thread order
 - Latency of SW scoreboard using Message Gateway used to be ~200 cycles
 - HW Scoreboard reduces it to a few cycles
- Two Scoreboard Types:
 - Stalling Scoreboard: a thread is stalled for dispatch until its scoreboard is cleared
 - Non-stalling Scoreboard: a thread can proceed until its first 'critical session' conditional load/store instruction SENDC



Thread ID and dependency

- Each thread is given an (x,y) ID
- Thread dependency is described by (x,y) delta distance
- A thread can have up to 8 dependences

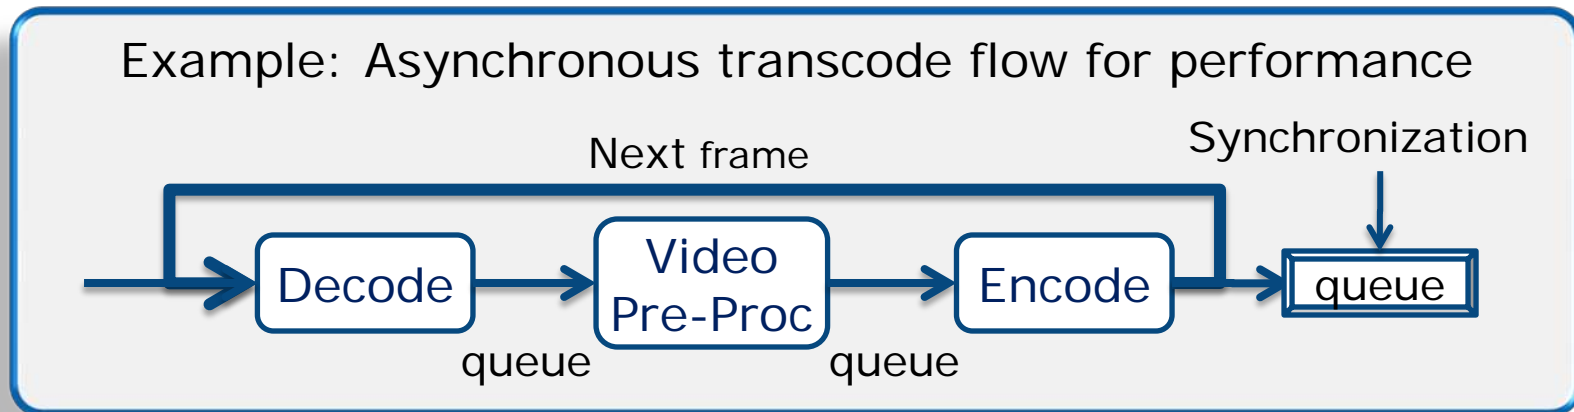
Video Processing – HW Acceleration

Media Sampler	Description
DN: De-Noise Filter	A motion detection based temporal (adaptive to noise history) & spatial (edge/texture preserved) de-noiser.
DI/FMD: De-interlacer & Film Mode Detection	A motion adaptive de-interlacer with multi-angle detection and interpolation, and with FMD statistics collection.
AVS: Advanced Video Scaler	A separable 8-tap poly-phase scaling filter with adaptive filter ringing suppression.
IEF: Image Enhancement Filter	A adaptive detail enhancement filter.
Media Pixel-OPs	Description
ACE: Adaptive Contrast Enhancement	Histogram equalization
STE: Skin-tune Enhancement	Skin color detection and enhancement
TCC: Total Color Control	Independent color adjustment with no effect on non-neighboring colors

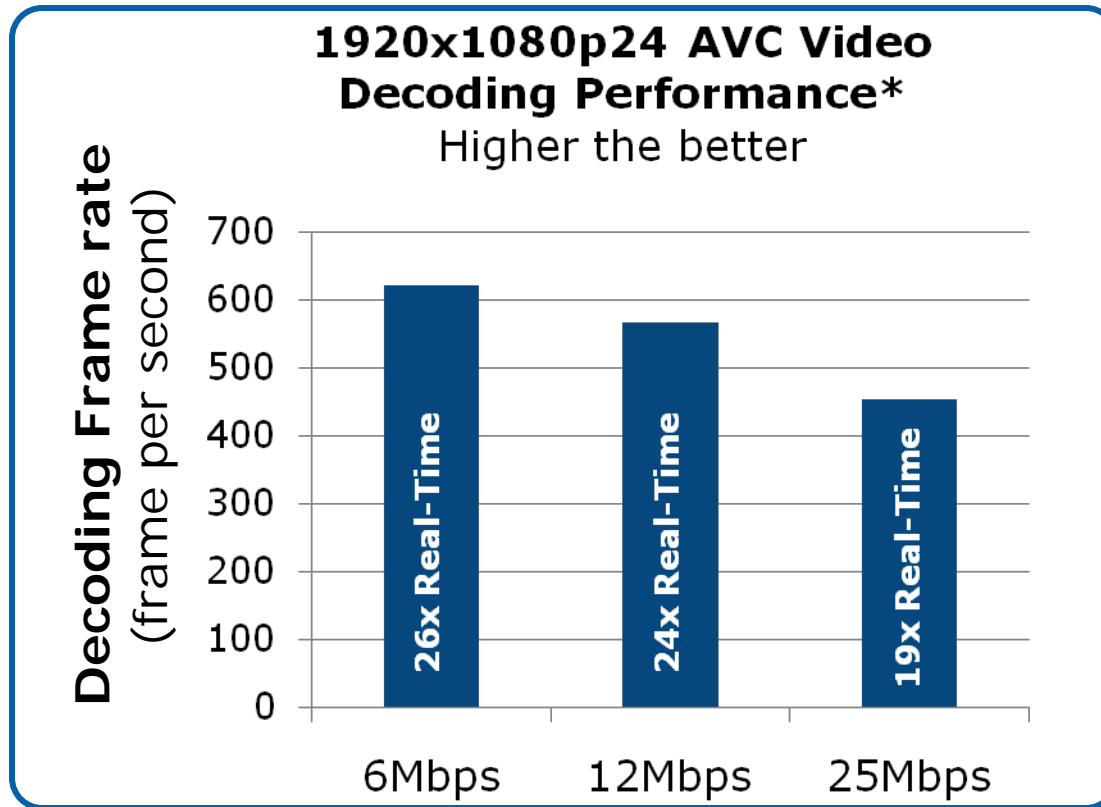
SW can form flexible video processing pipeline for either playback or video transcoding

High Performance Video Transcode

- PG provides HW acceleration for all stages of video transcode
- Video transcode is enabled through Intel® Media SDK
 - Necessary software stack to expose PG's Video Pre-processing and Encoding functions
 - Simplified SW stack for PG's Video Decoding functions (w.r.t DXVA)
 - Quality and performance modes provided for different usage cases



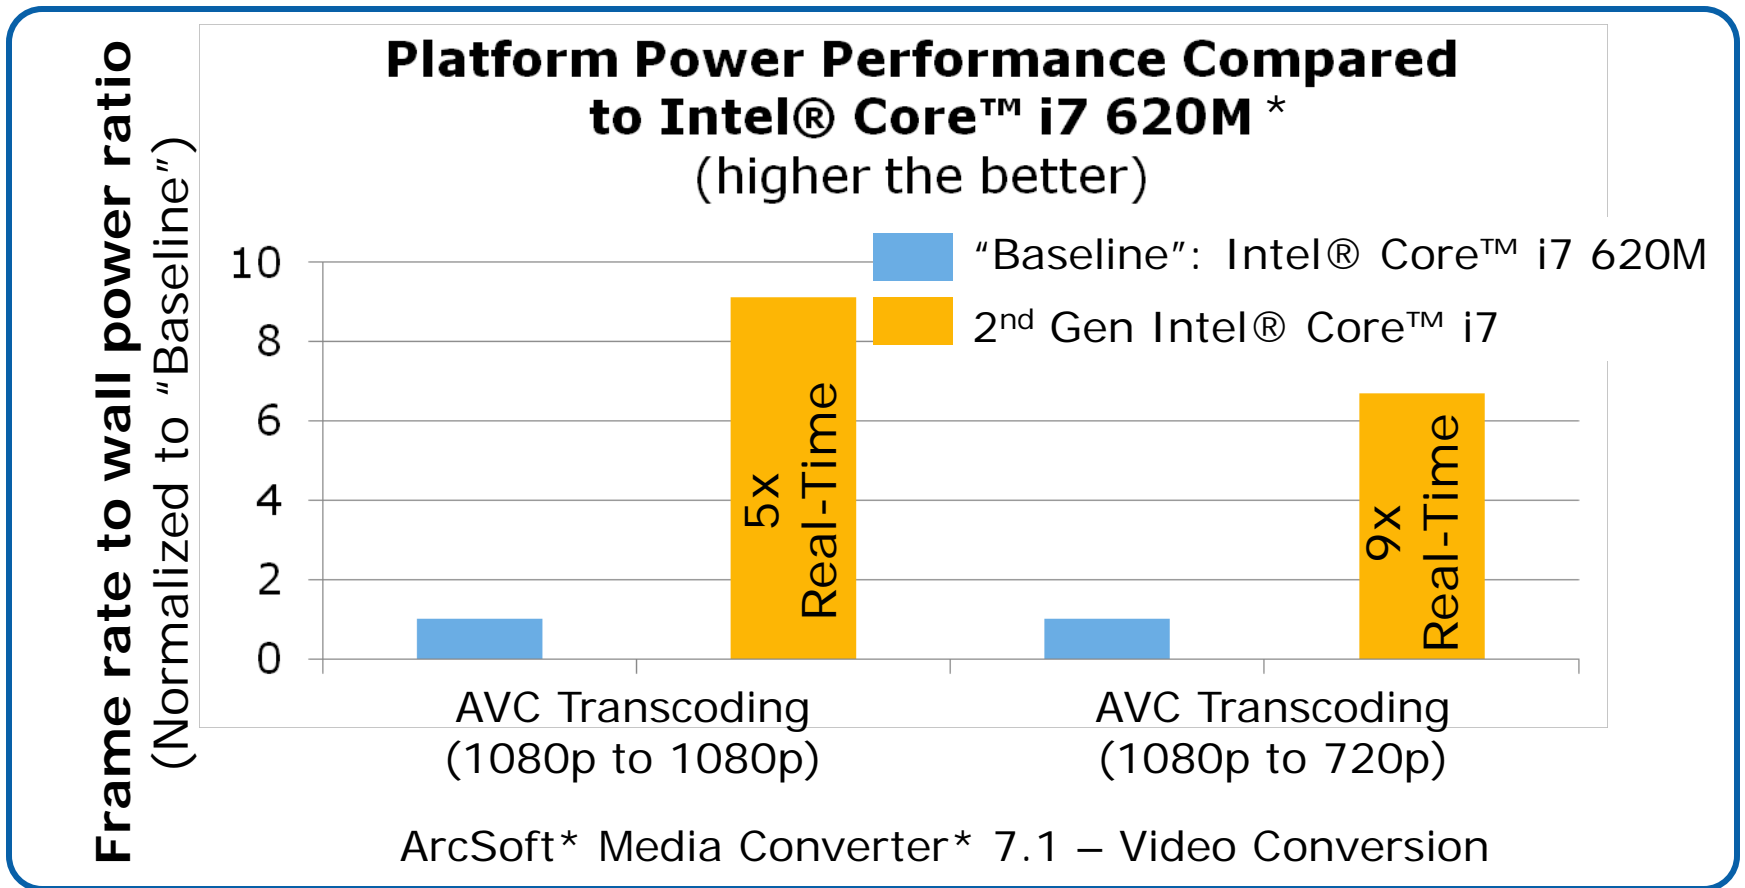
Video Decode Performance



- Potential usages enabled by high speed video decode
 - Fast video transcode
 - Video editing: trick mode, fast rewind, live user interaction with video
 - Video Wall: multiple videos on screen for content preview or security

* See configuration details in Appendix

Intel® Quick Sync Video Power Efficiency



- Quick Sync Video provides HD transcoding much faster than real-time
- Power efficiency is 6-9x over previous gen Intel Core™ processor

* See configuration details in Appendix

Concluding Remarks

- Intel® Quick Sync Video is a unique video transcoding implementation on Processor Graphics
 - The PG-based programmable solution, assisted with HW accelerators, delivers on the promise of performance, quality and power
- The underlining techniques also enable broad applications, for example
 - Intel® Intra 3D Technology: MFX decoder has ample headroom for 2X HD decoding for Blu-Ray Stereoscopic 3D playback
 - Intel® Wireless Display: Quick Sync Video enables real-time 1080p encoding for Wireless Display

Backups



Relevant Web Links

- Intel® Open Source HD Graphics Programmer's Reference Manual (PRM)
 - <http://intellinuxgraphics.org/index.html>
- Intel® Media SDK
 - <http://www.intel.com/software/mediasdk>
- MPEG-4 AVC/H.264 Video Codecs Comparison, MSU, May 2011
 - http://www.compression.ru/video/codec_comparison/h264_2011/mpeg-4_avc_h264_video_codecs_comparison.pdf

Appendix

- Source: Intel
- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.
- ArcSoft* Media Converter* 7.1 - Video conversion workload consist of a 4 minutes 40 seconds, 25mbps, 1920x1080p, AVC/H.264 benchmark video file. The file is transcoded to a lower bit rate, 12mbps, 1920x1080p, AVC file and a lower bit rate, 10mbps, lower resolution, 1280x720p, AVC file.
- 2nd Gen Intel® Core™ i7: Mobile Core i7-2820QM (2.4GHz, Turbo 3.5GHz, GT 650MHz, Turbo 1.3GHz), DDR3-1600 2CH 2GB, Windows 7 32-bit with PV2 driver (Build 2246)
- Intel® Core™ i7 620M: Core i7 620M (2.6GHz, Turbo 3.3GHz), DDR3-1333MHz 2CH 1GB.

