PHILIPS

HOTCHIPS 2006

Heterogeneous multiprocessing for efficient multi-standard high definition video decoding

Outline

- Multi-Standard Video Decoder overview
  - Internal system architecture
- Design challenges
  - CPU performance
  - Customized TriMedia VLIW core
  - Programmable pixel processing
- Results & Conclusion

MSVD charter

- Integrated video decoding solution for consumer application
  - Support for common video compression formats in the consumer world
  - Up to high definition resolution
  - Concurrent stream decoding
  - Supporting consumer encoded ‘ almost compliant ’ material
  - Advanced trick modes
  - Advanced post processing for best in class picture quality
  - Competitive against fully hardwired solutions

PHILIPS

MSVD system feature list (1)

- Supported formats
  - H.264 (aka MPEG4 AVC or MPEG4 part 10) High Profile @ Level4 (1920x1080 60Hz interlaced or 1080i, 1280x720 60Hz progressive or 720p)
  - VC-1 (WM9) Advanced Profile @ Level 3 (1080i, 720p)
  - MPEG4 Advanced Simple Profile @ Level3 (1080i, 720p)
  - DivX 3.11/ 4.x.x / 5.x.x / 6 HD profile
  - MPEG1/2 Main Profile @ High Level (1080i, 720p)
  - DV / DVCAM / DVCPRO-25 @ 25 Mbits/s
  - JPEG EXIF2.2 format up to 8192x8192
- Video post-processing
  - Deblocking
  - Deringing

MSVD system overview

- Control / parsing in firmware
  - Running on embedded controller
  - low amount of control glue
- HW assistance for stream parsing
  - Performance independent of bit rate
- Autonomous pixel crunching
  - Using a set of dedicated units
  - HD level performance at consumer cost

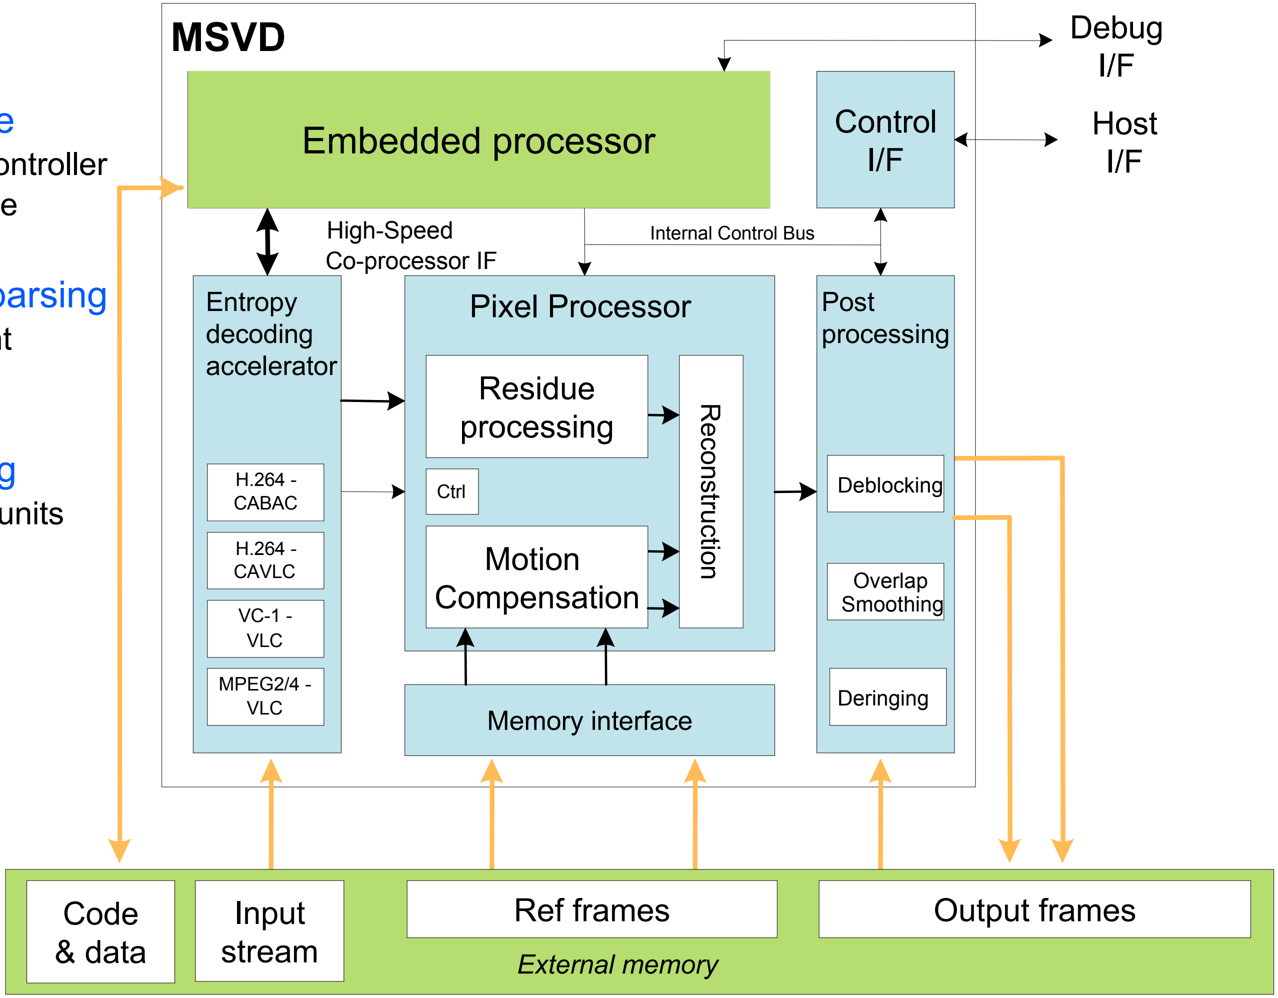

MSVD system overview

- Benefits
  - Good range of standard decoders in limited Si area
  - High decoding throughput
  - Easier support for non compliant streams (PC world)
  - Field upgradeable firmware
  - Possibility to tweak error concealment / playability
  - Possibility to expand supported formats
    - If using similar pixel crunching functions as already supported codecs (e.g. MPEG4 variants: DivX/Xvid)
  - Scalable
    - Reduced set of standard
    - Low operating frequency for low power standard definition decoding
    - Choice of embedded controller

Outline

- Multi-Standard Video Decoder overview
  - Internal system architecture
- Design challenges
  - CPU performance
  - Customized TriMedia VLIW core
  - Programmable pixel processing
- Results & Conclusion

CPU performance - Requirements (1)

- Origin of coding efficiency improvements
  - Deeper picture segmentation
    - Down to 4x4 blocks
  - Sophisticated prediction scheme
    - Most symbols are predicted
    - Adaptive selection of predictor for each symbol
  - Advanced entropy coding techniques
    - Arithmetic coding / Multiplication of Huffman tables (VC-1)
    - Most symbols are now predicted and not simply transmitted in bit stream (I.e. MB type/CBP/QP)

- Consequence
  - More symbols to decode
  - Steep increase in computation and data manipulation required for decoding a symbol
  - Example motion vectors prediction
    - 1 load / 1 store per MV, up to 4 MVs per MB in MPEG2
    - 12 load / 3 store / 3 compare per MV, up to 32 MV's per MB in H.264

PHILIPS

CPU performance - Requirements (2)

- High impact on stream parsing complexity
  - 8x more performance required from MPEG2 to H.264
- Parsing Performance requirements for 1080i
  - MB rate: 245760 MB/s
    - 1354 cycles / MB @ 333MHz
      - 6 cycles/symbol (incl. 384/2 transform coefficient)
    - Complexity (excluding transform coefficient parsing)
      - 700 load / store operations / MB for context manipulation
      - Total of 3000 ops / MB
      - High level of execution hazards (100 branches)

PHILIPS

CPU performance - TM Config

- Configurable TriMedia processor
  - Based on TM3270 architecture
    - 5 issue slot VLIW architecture
    - 35 functional units
    - 9 stage pipeline
  - Configurable
    - Instruction cache: 32kB
    - Data cache: 16kB
    - Register file size: 96 General Purpose Registers
    - Dedicated coprocessor interface
      - Compiler supports optimal scheduling of accesses
    - Customizable function unit: SIMD instructions
      - Dual operation for MV handling (dual add/median)
  - Parallel LOAD and STORE unit
    - 1 load and 1 store per cycle
    - Increase parallelism by reducing pressure on the load / store unit

VLIW TriMedia Configurable Core

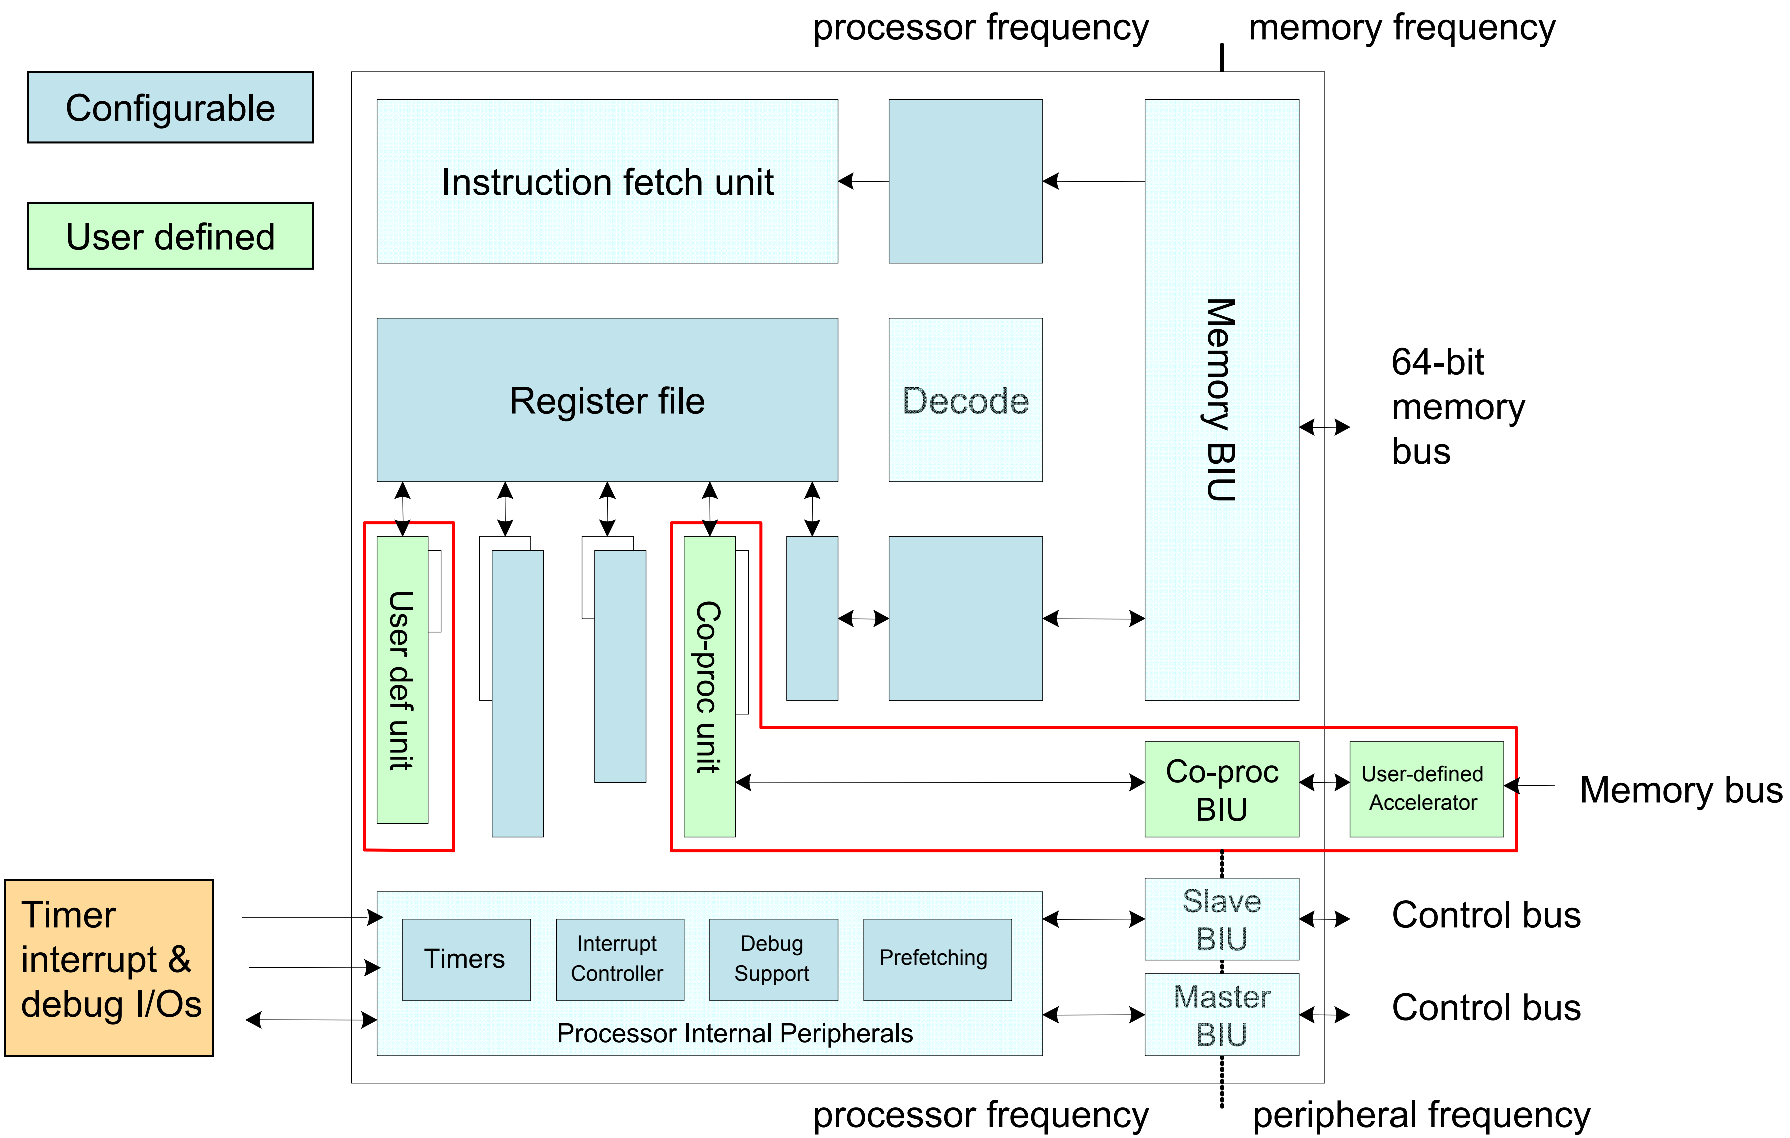

TM Config - Coprocessor BIU

- Connected to Entropy Decoding Accelerator
  - Primitives :
    - rd_bits(n) / rd_uvlc(n) / rd_cabac_symbol(ctx0,ctx1,…)
- Load/Store programmer's interface:
  - cop_ld32r, cop_ld32r, cop_ld32x, cop_st32d
- Configurable coprocessor load / store latency
  - Latency set to 7 cycles for LD / 4 cycles for ST
  - Scheduling support in compiler tool chain for 0 overhead access
- Drastically reduces stalls on co-processor accesses
  - 15 % performance gain compared to normal BIU

TM Config – User defined unit

- Customizable SIMD unit in CPU pipeline
  - Up to 4 32-bit sources / Up to 2 32-bit destinations
  - Full compiler support
  - RTL module inserted at synthesis time
- Examples of usage
  - H264 motion prediction
  - Motion vector scaling
  - H264 delta motion vector context computation
- Benefits
  - Parallelism of SIMD (e.g. (X,Y) computed in single step)
  - Removal of low level branches
    - Very damaging for efficient scheduling on VLIW with long pipeline

Pixel processing challenges

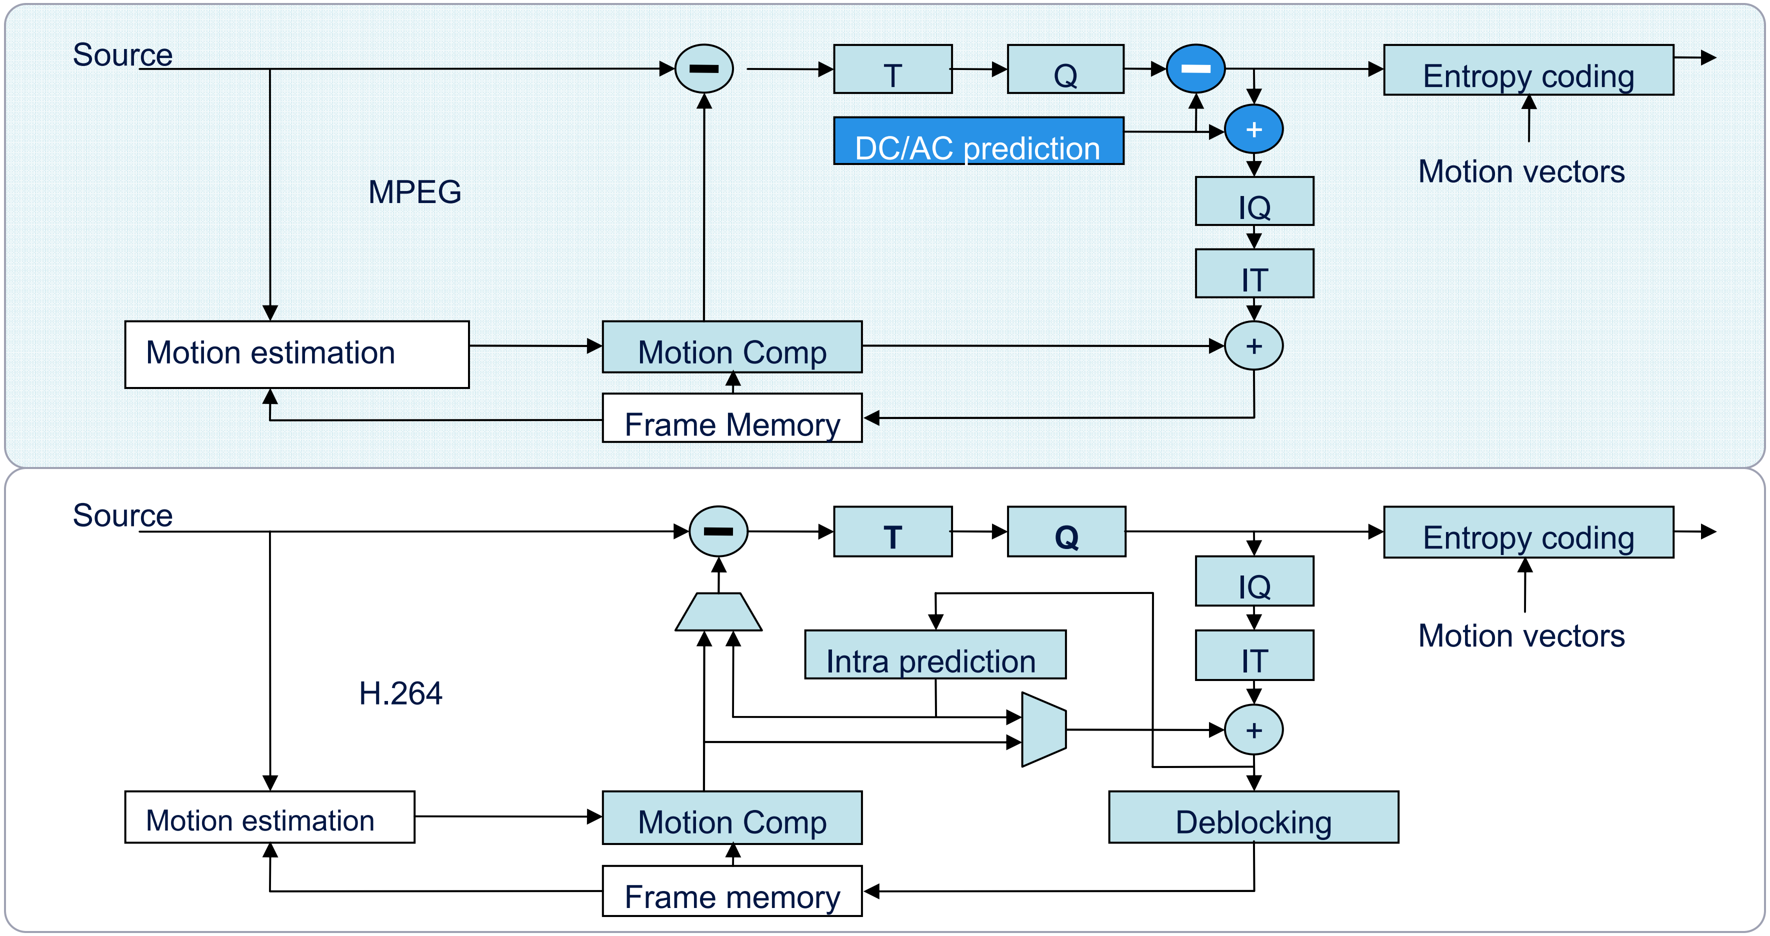

Pixel processing challenges

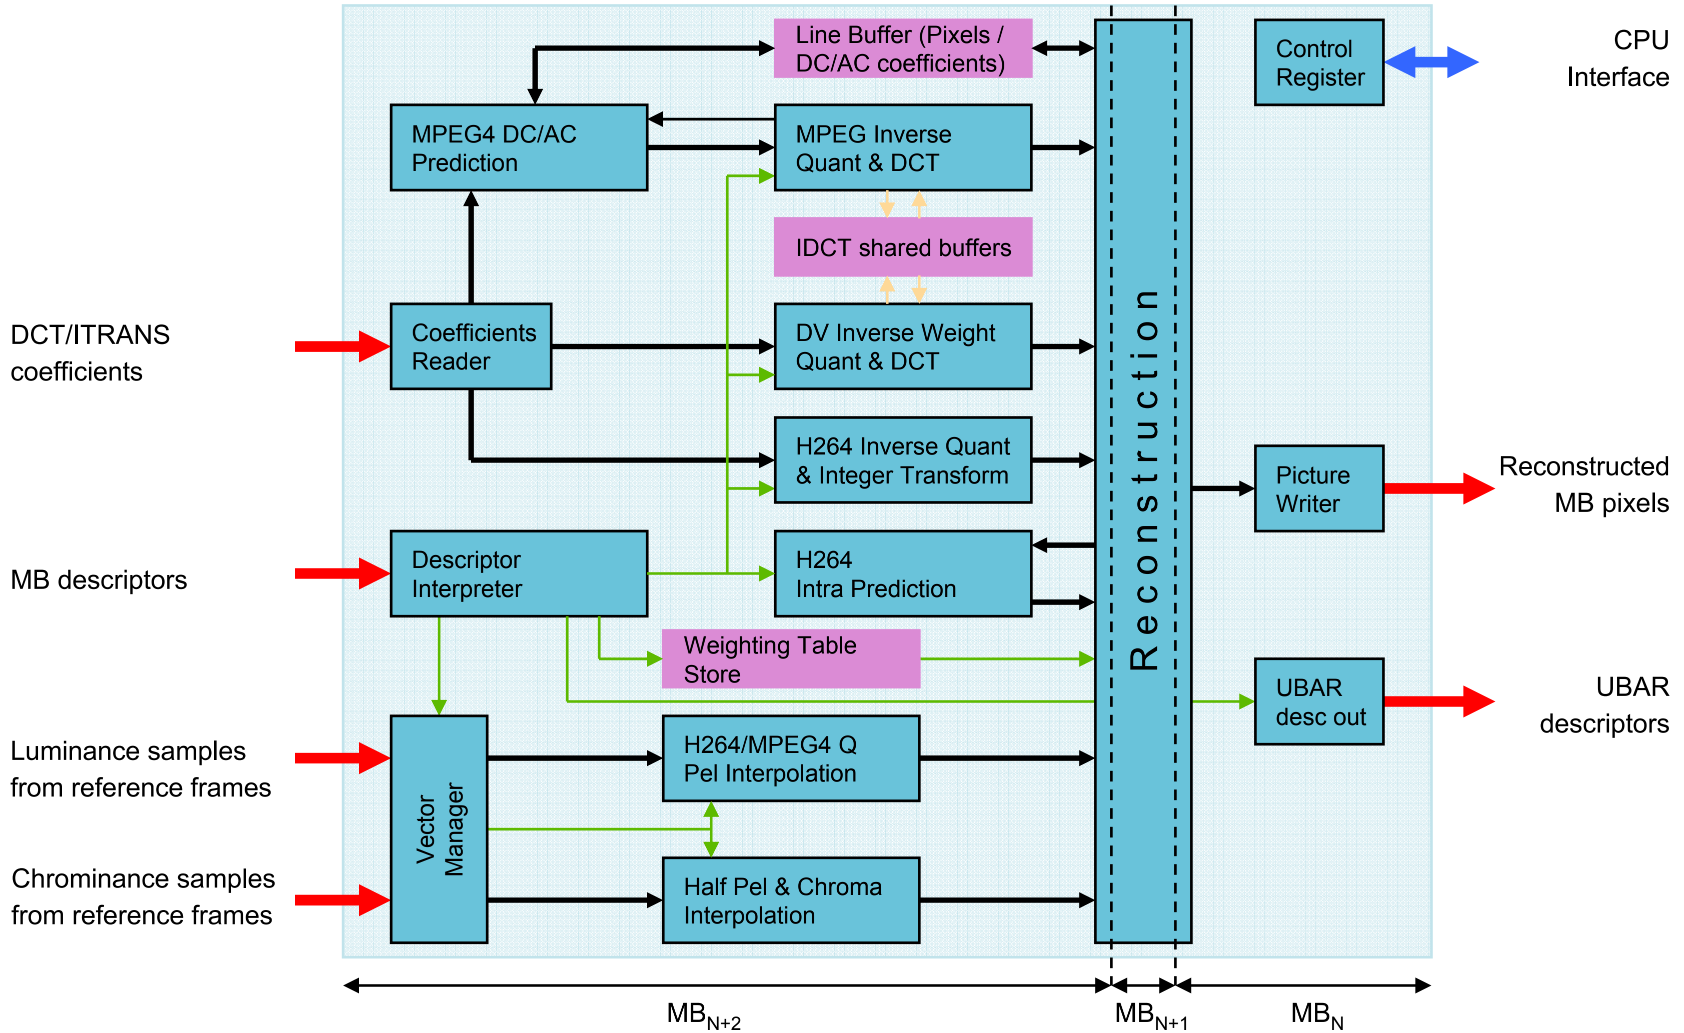

Programmable pixel processors

- Same principles used for video compression
  - Transform based residue / intra coding
  - Quantization
  - Combination of motion prediction and residue
  - Filtering for motion predictors
  - Deblocking
- But different processes and implementation
- Consequence for a multi standard solution
  - Hardwired solutions become complex
    - Resource sharing is possible but increase verification effort
  - Symmetric multi processing is not efficient
    - Diversity of algorithms defeats architecture optimization
    - Scheduling is an issue due to increasing data dependency

Programmable pixel processors

- Example: Transform
  - MPEG started with 8x8 IDCT but over time
    - Different shapes
    - Different dynamic range / rounding
    - Integer transforms
  - Complexity
    - 23.6 M 1D transform / s for MPEG2 HD
    - Operations for 1D DCT
      - 50 operations including 8 input loads + 8 output writes
    - Budget: 8 cycles at 200 MHz
  - Need for a dedicated processor structure removing overhead of typical architectures
    - Direct access to data is key to achieve performance and efficiency

Programmable pixel processors

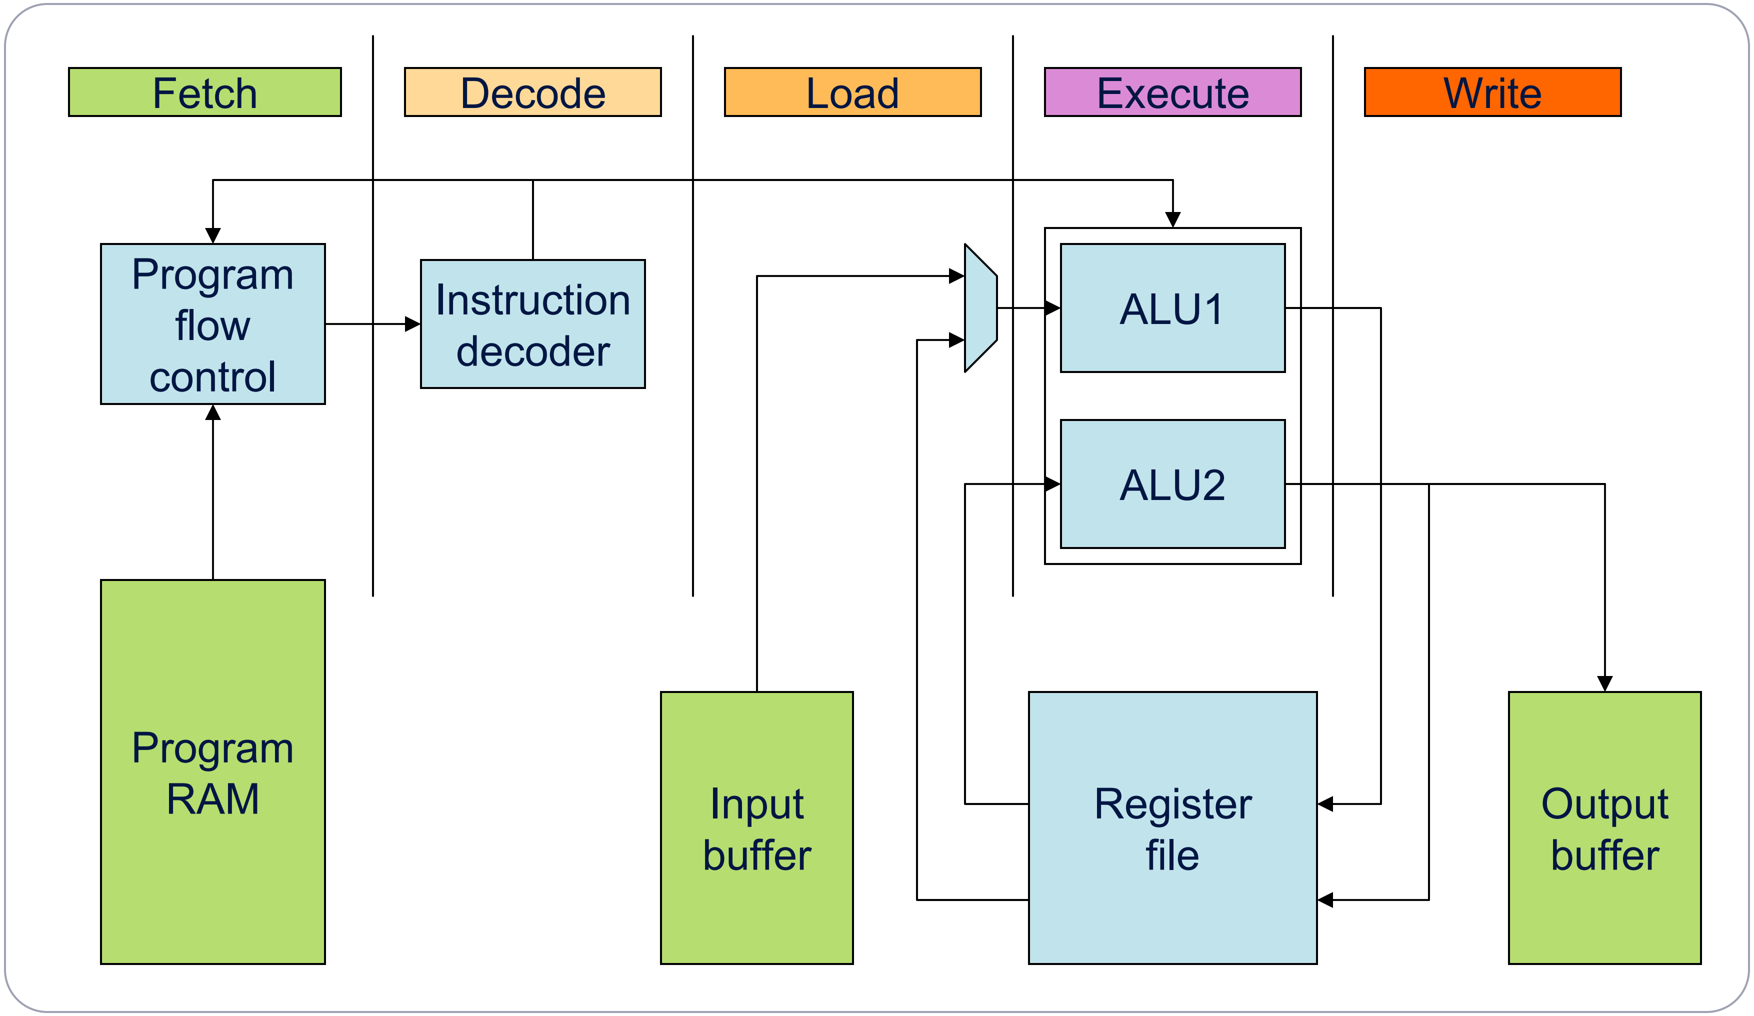

Programmable pixel processors

- Solution
  - Dual issue 5 stages ASIP core
  - Direct input / output buffer access
  - 24 general purpose registers
  - 2 ALUs
    - Single cycle butterfly op
    - Load / butterfly op
    - Butterfly, round and store op
  - Zero overhead loop support
  - Different transforms supported by
    - Adjusting coefficients in GPR for butterfly ops
    - Changing loop limits (shapes)
    - Adjusting rounding parameters (in dedicated registers)
- Performance
  - MPEG DCT in 112 cycles (7 cycle / 1D IDCT)
  - H.264 4x4 transform in 16 cycles

Programmable pixel processors

- Results
  - Hardwired solution for 5 standards : 0.2 mm2 in 90 nm LP
  - Programmable engine : 0.12 mm2 in 90 nm LP

- Other application : deblocking
  - VC-1 has 2 deblocking mechanisms, overlap transform and deblocking filter
  - Control in hardware is possible but difficult to mature
  - New unit is implemented as a programmable engine with Lisatek tool suite from CoWare
    - Pipelined filter operation
    - Control moved to firmware
    - Improvement in debug time for no area penalty
      - Area improves as other deblocking algorithms are mapped onto the same processor

Programmable pixel processors

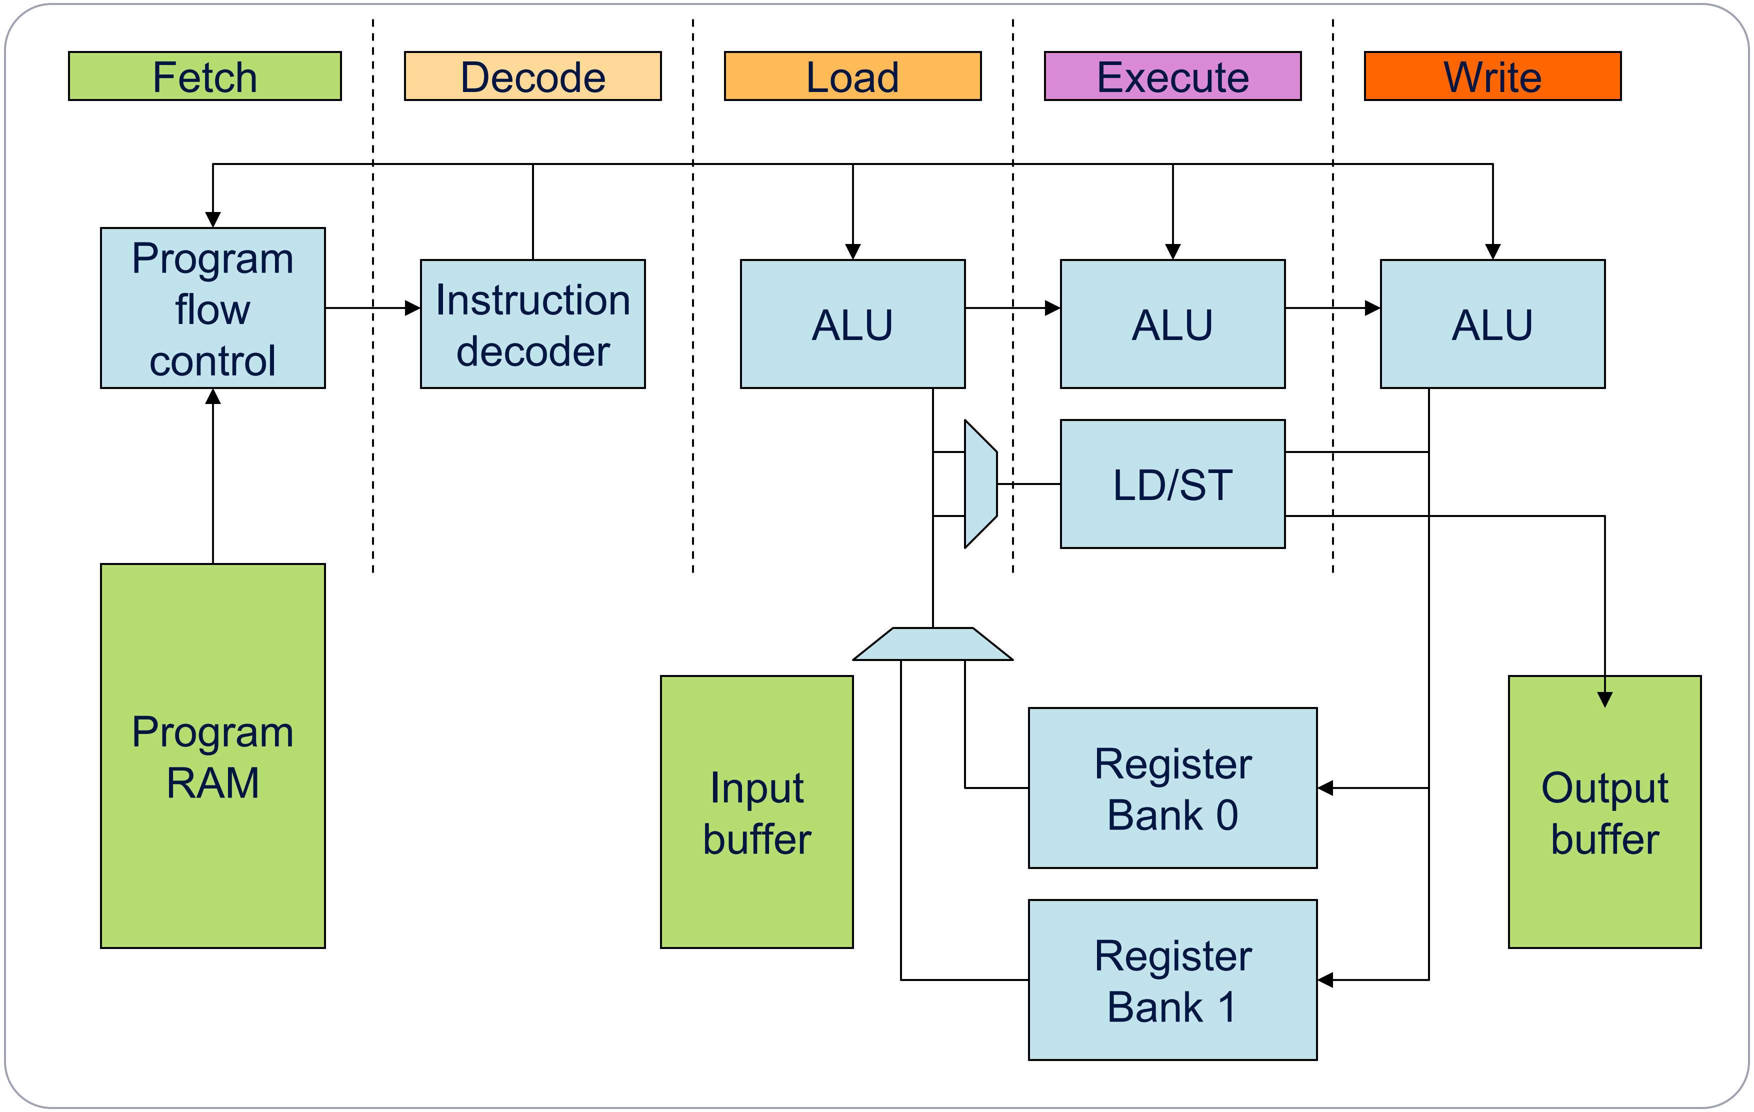

Outline

- Multi-Standard Video Decoder overview
  - Internal system architecture
- Design challenges
  - CPU performance
  - Customized TriMedia VLIW core
  - Programmable pixel processing
- Results & Conclusion

Results - MSVD key figures

- Hardware characteristics
  - All figures are in 90 nm low power process
  - CPU core operating @ 333 MHz (soft core)
    - 5 issues VLIW architecture + support for SIMD
    - 32/16 kB I/D cache
    - Customized instructions and 0 overhead coprocessor
    - Area 3.2 mm2 including caches
  - MSVD core operating @ 166 MHz
    - Area 3.70 mm^2
  - Total area 6.9 mm^2

Conclusion

- Efficient multi-standard HD solution achieved by
  - Proper partitioning of tasks
  - Usage of customized CPU core associated with closely coupled and loosely coupled dedicated processing units
- Preserve a good deal of flexibility at HD level performance
  - Wide range of standards supported
  - Capability to deal with variants (e.g. DivX x.xx)
  - Error concealment strategy in SW
- Programmable computation engines bring flexibility, design time and area reduction

PHILIPS