SoCs for Portable Video Applications: Architecture level Considerations

Mahesh Mehendale
m-mehendale@ti.com

EDP 2007

Agenda

- Video processing requirements of portable entertainment applications
- Characterizing “variability” in digital video processing
- Low power design techniques and their applicability in the context of Digital Video Sub-system and the SoC
- EDA challenges and Opportunities

Video Processing Chain

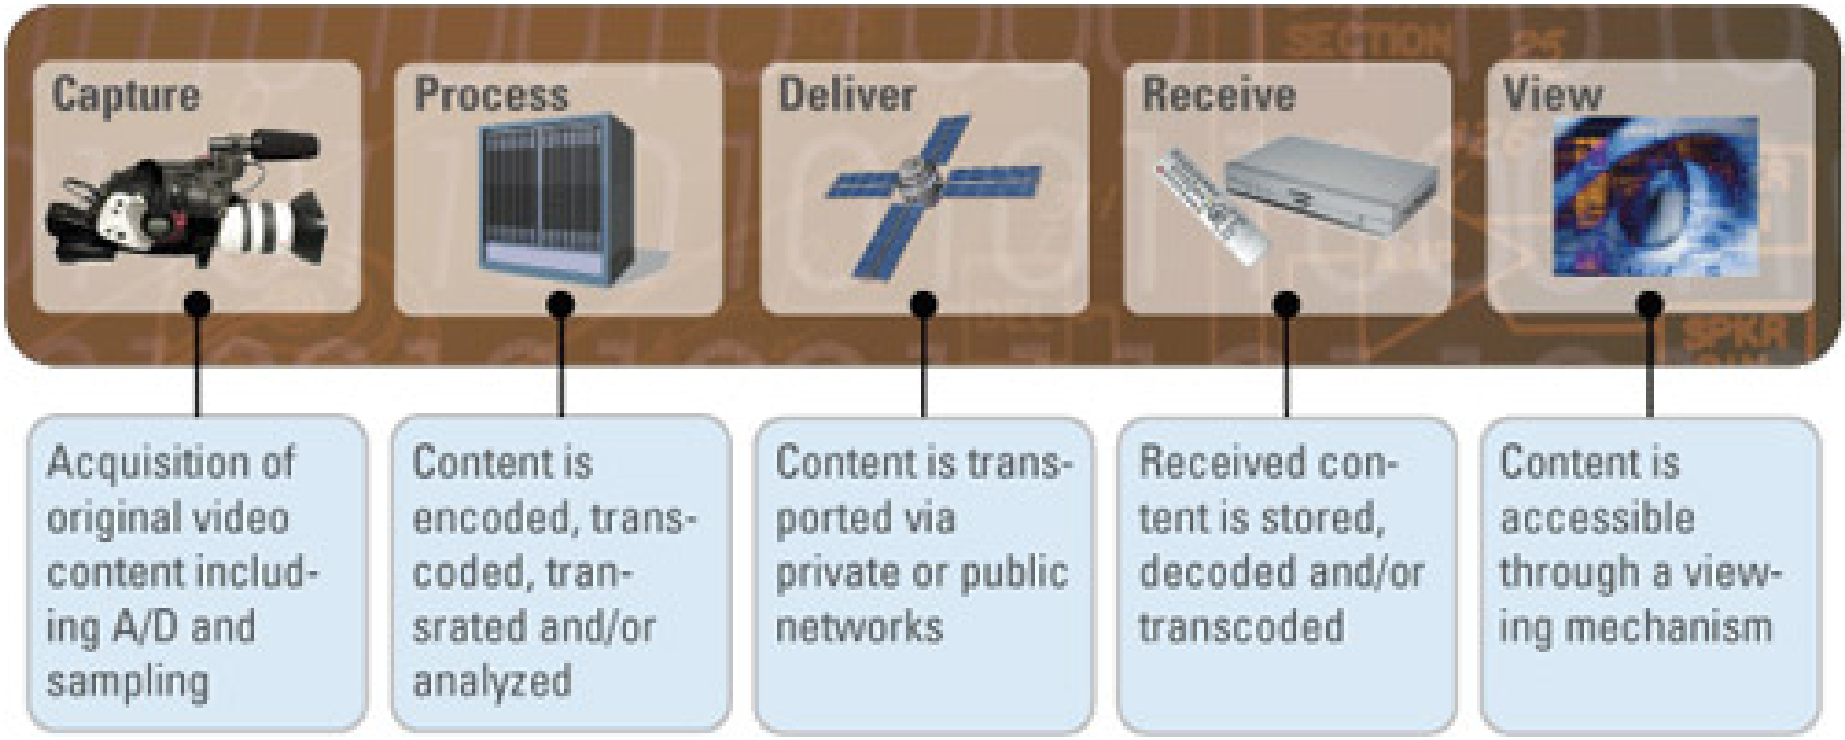

Personal Video Entertainment

- Portable Video Recorder
- Portable TV (DVB-T, DVB-H)
- Portable Media Player
- Digital Camcorder
- Portable Navigation
- Video phone
- Web terminal

Portable Media Player – video interfaces

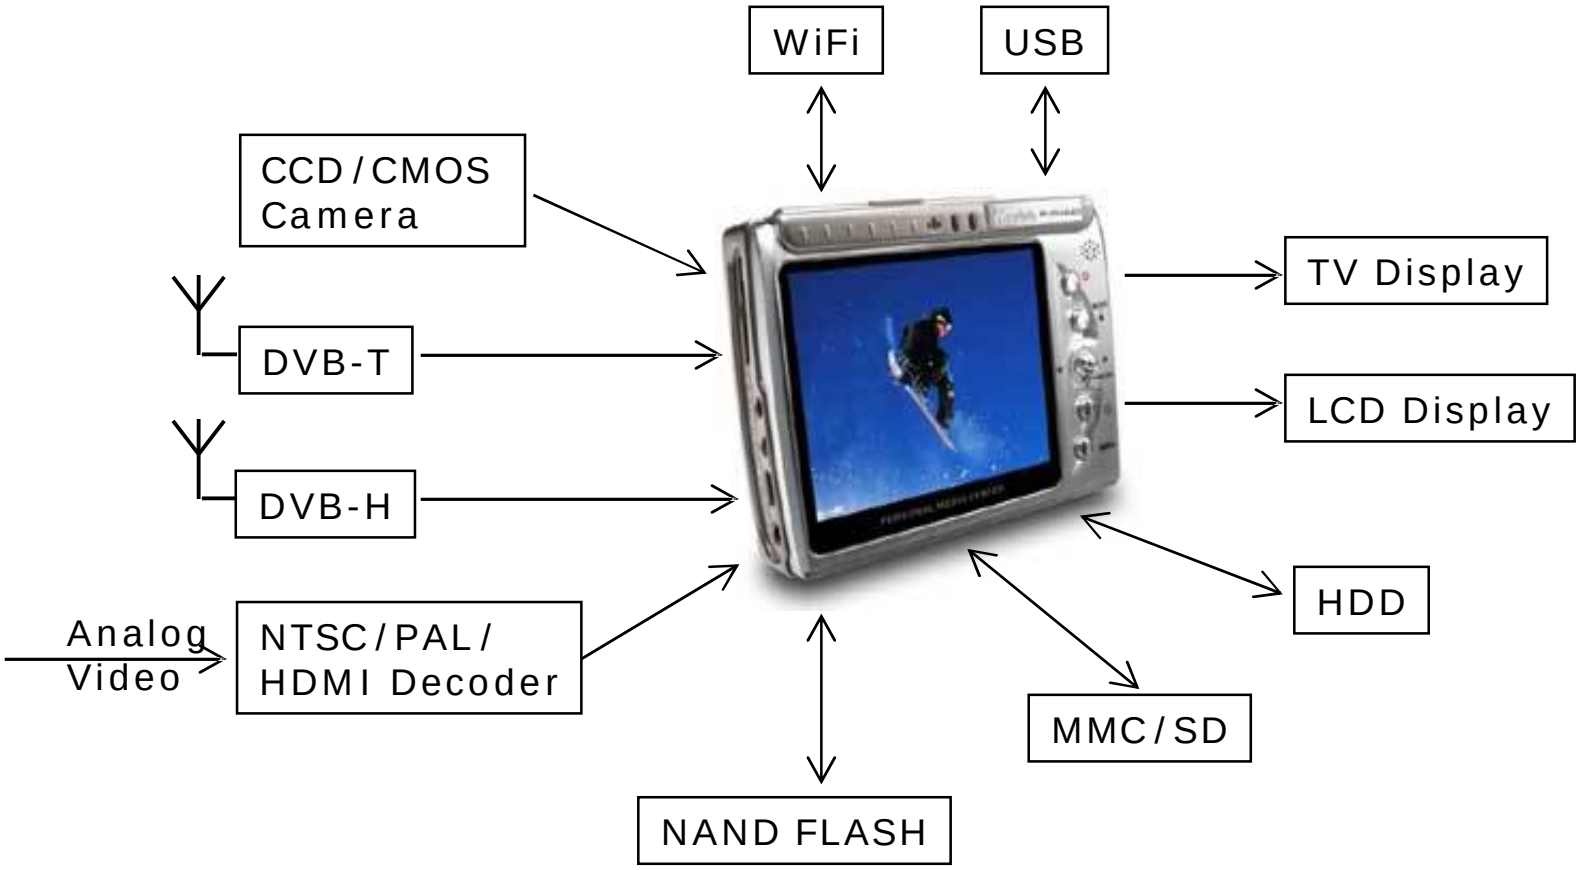

Customer care-abouts

- Multi-standard, multi-format video processing

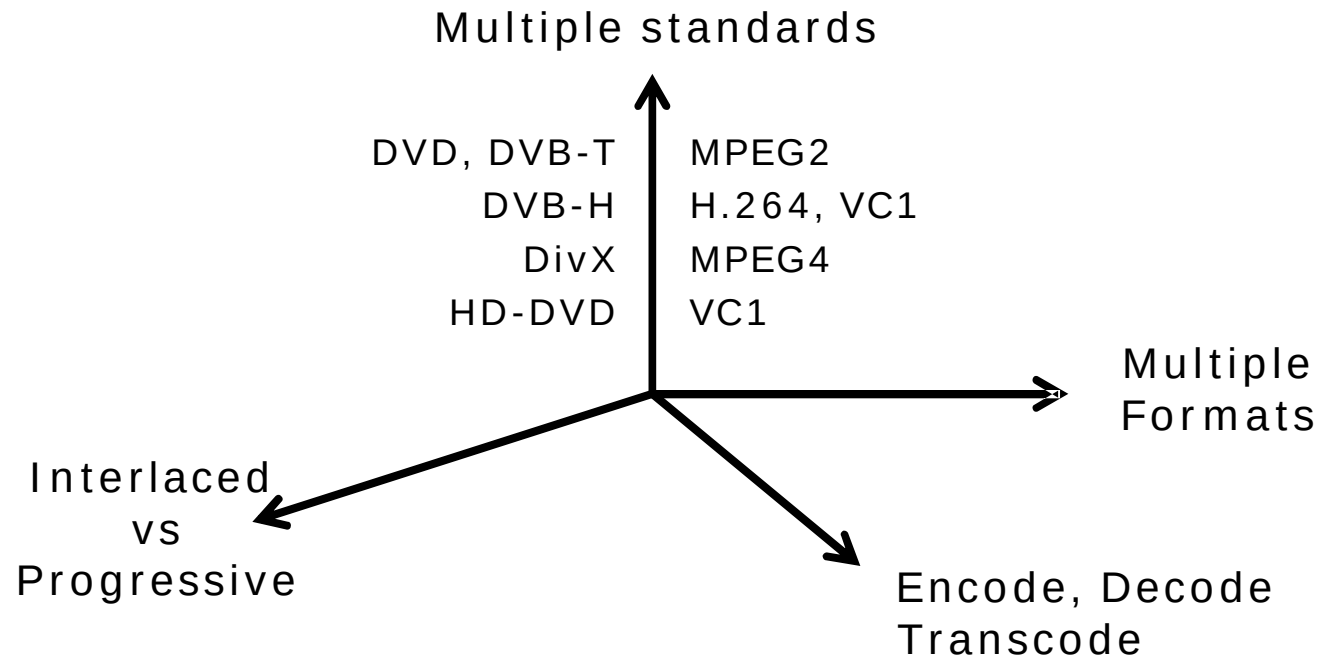

- Cost
- Power/Energy

Video formats

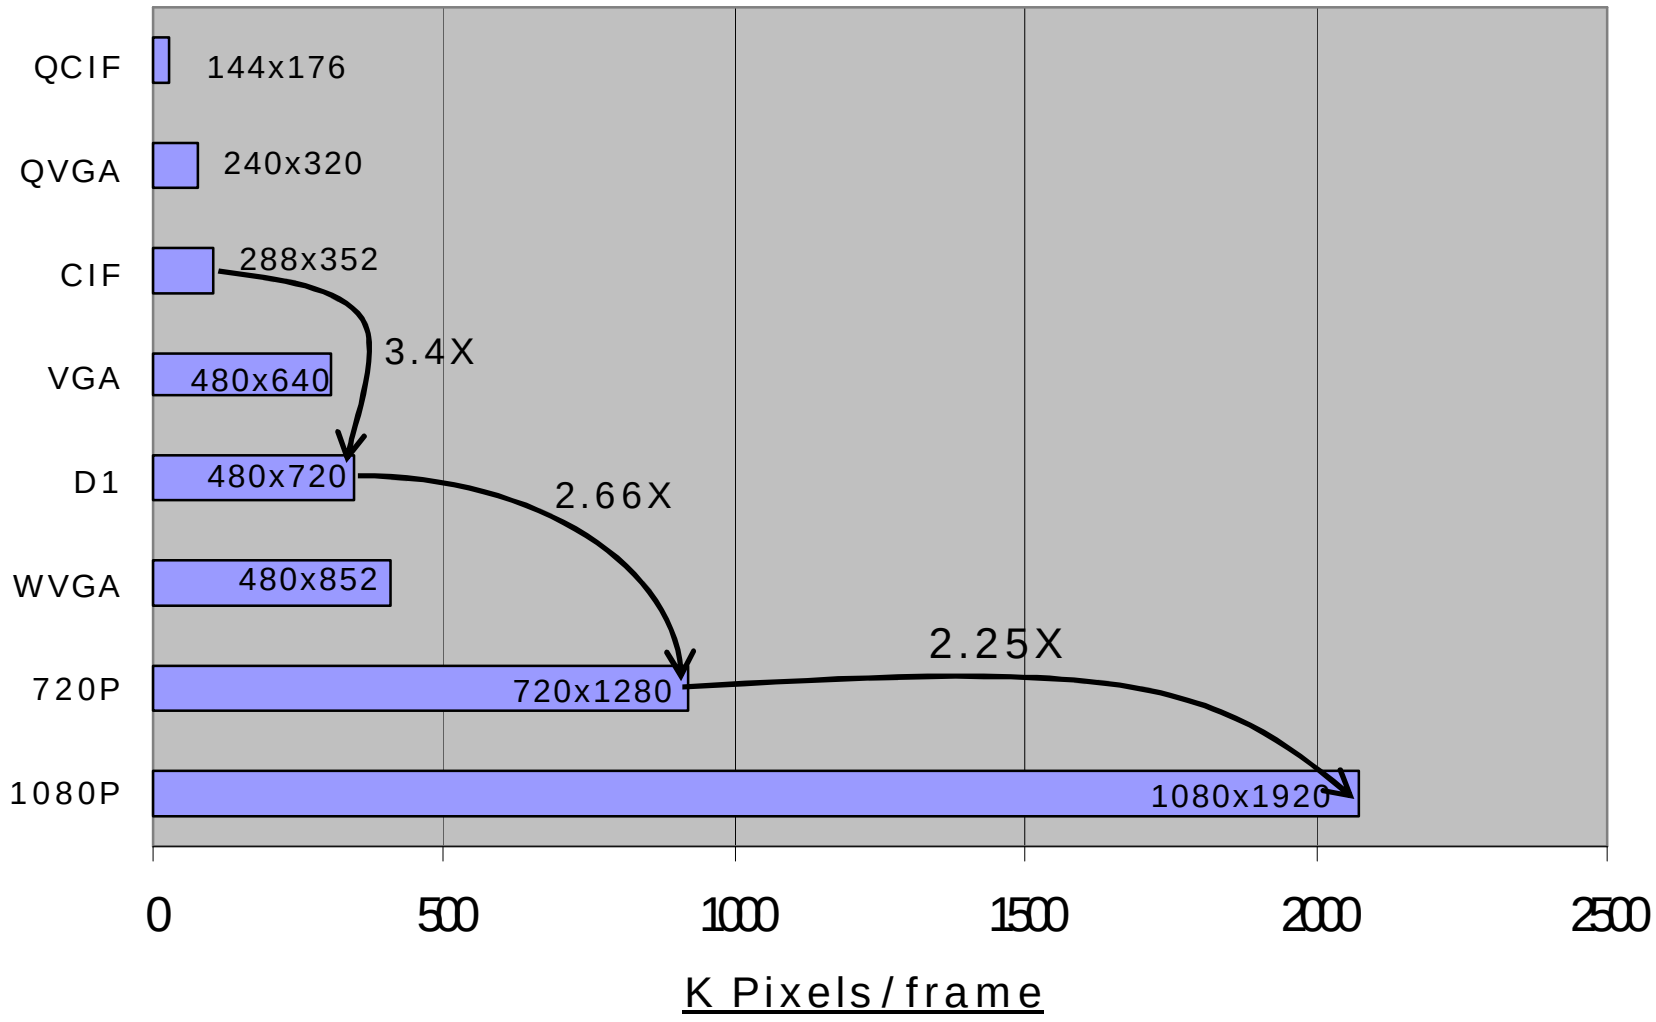

DV Engine Solution Space

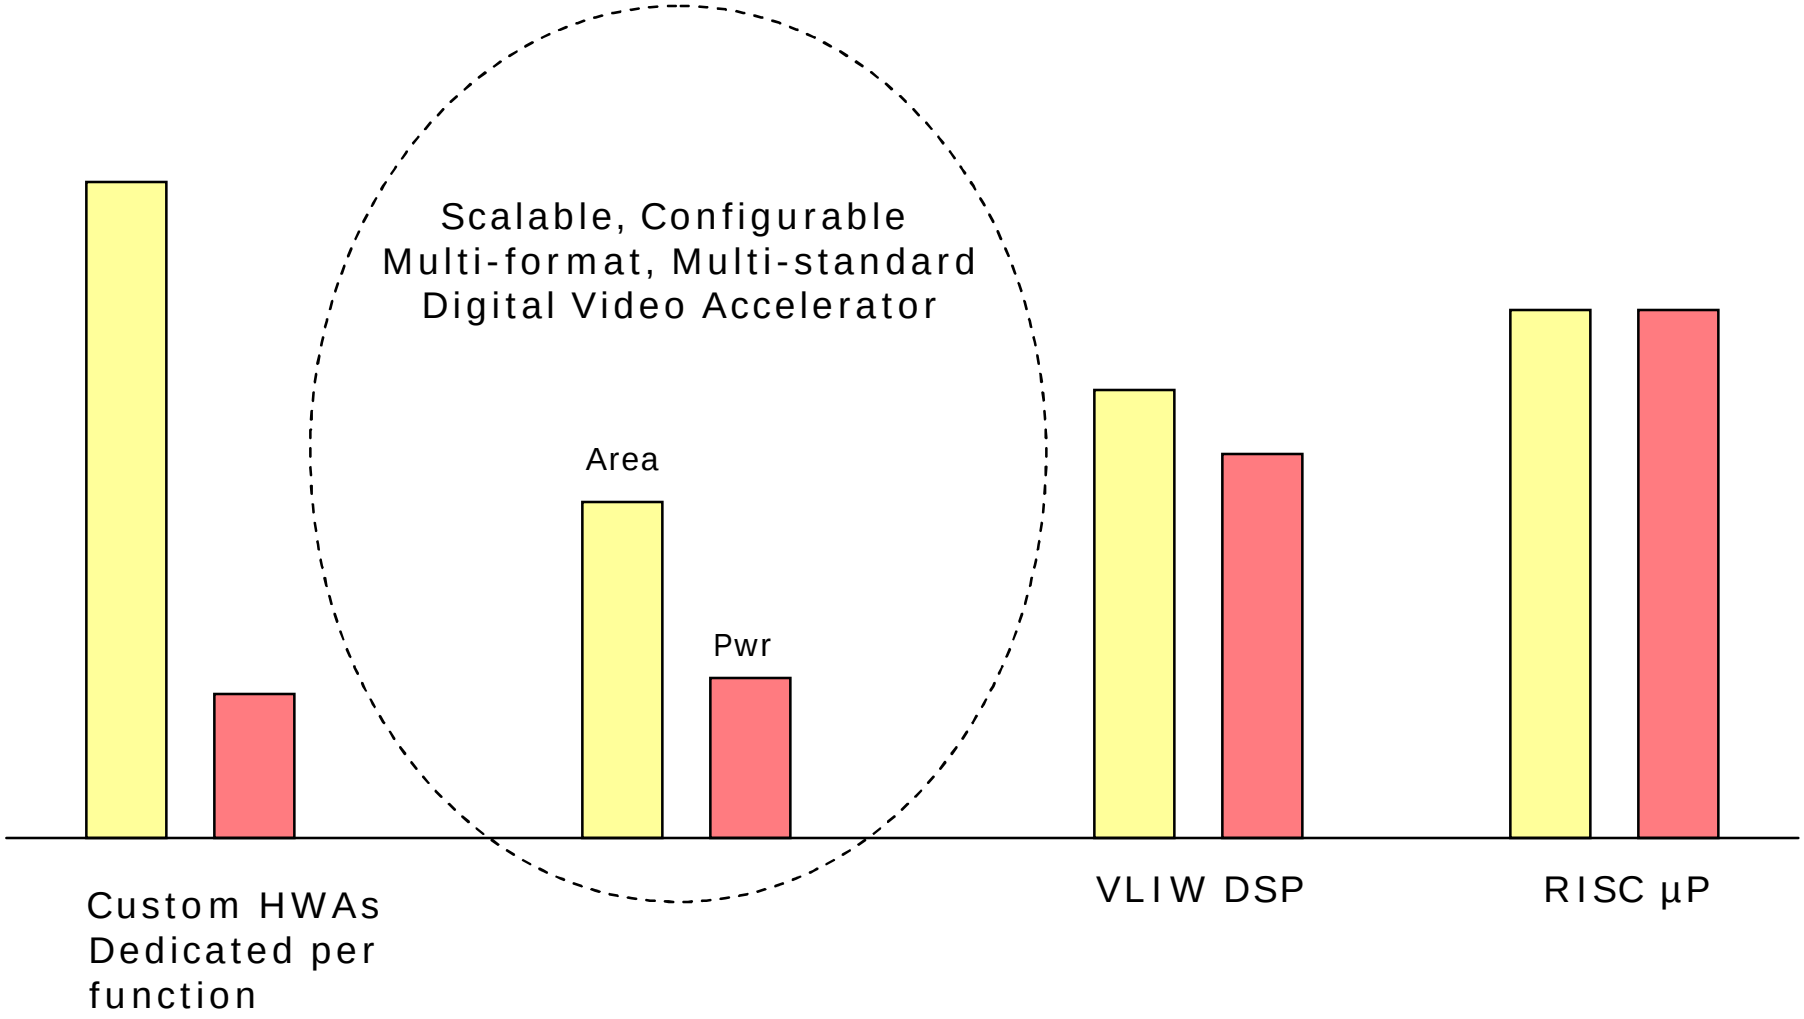

H.264 Encoder

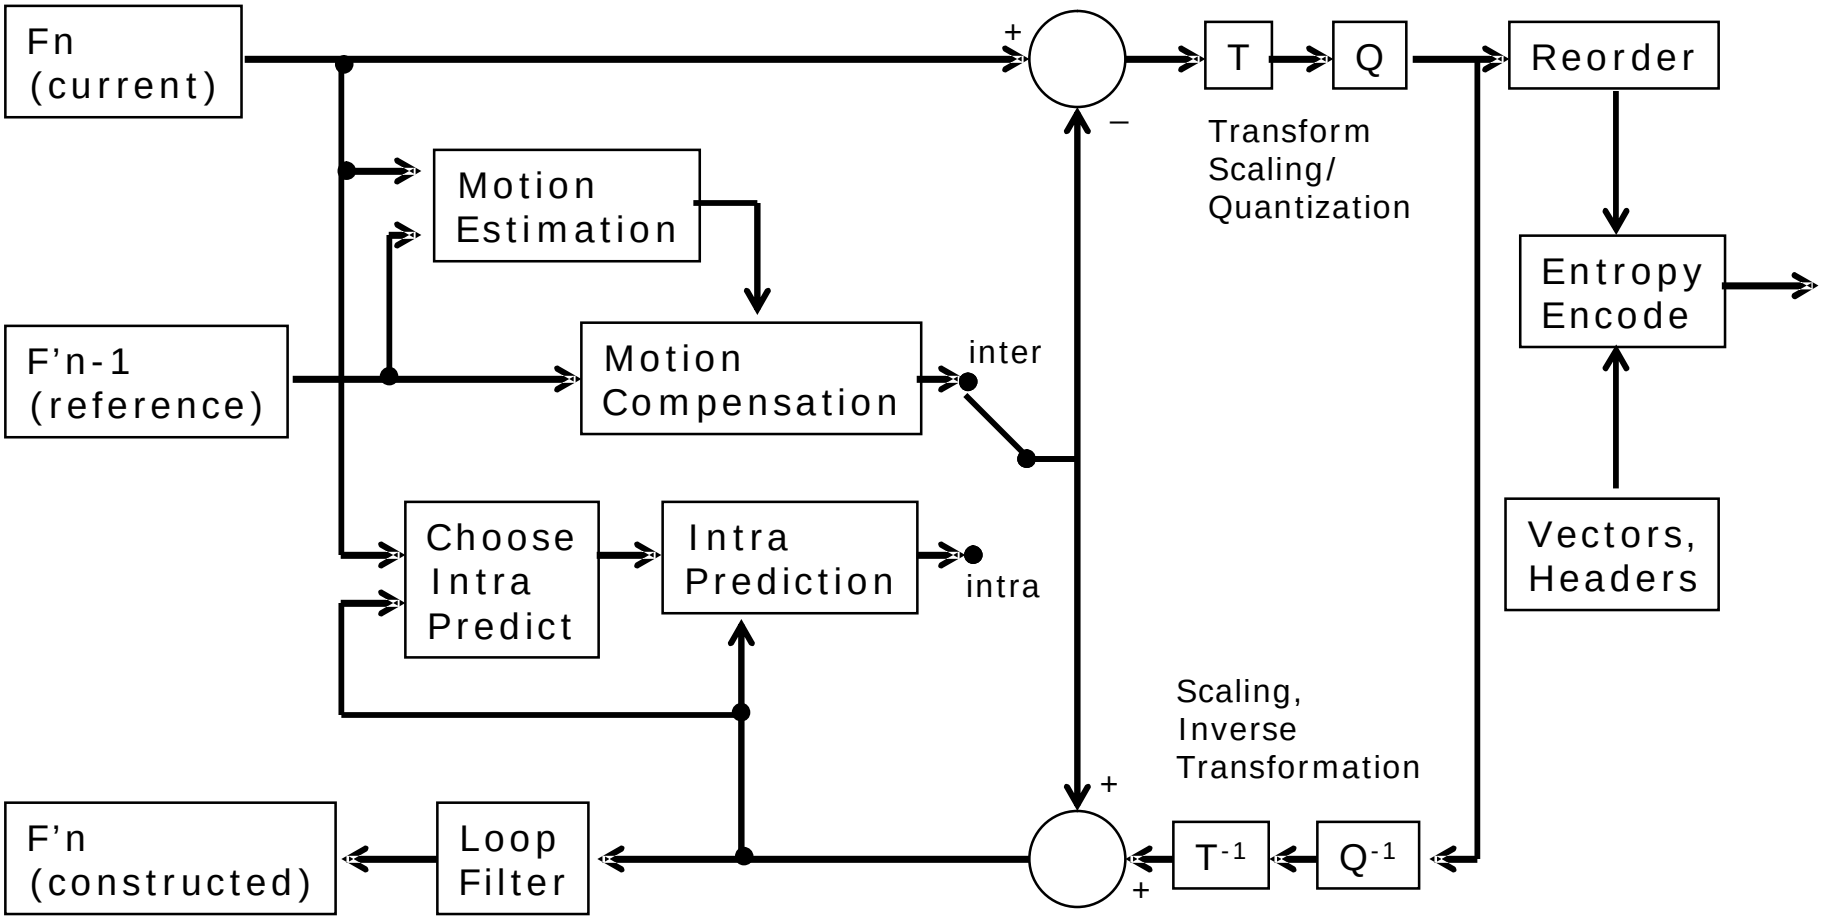

H.264 Decoder

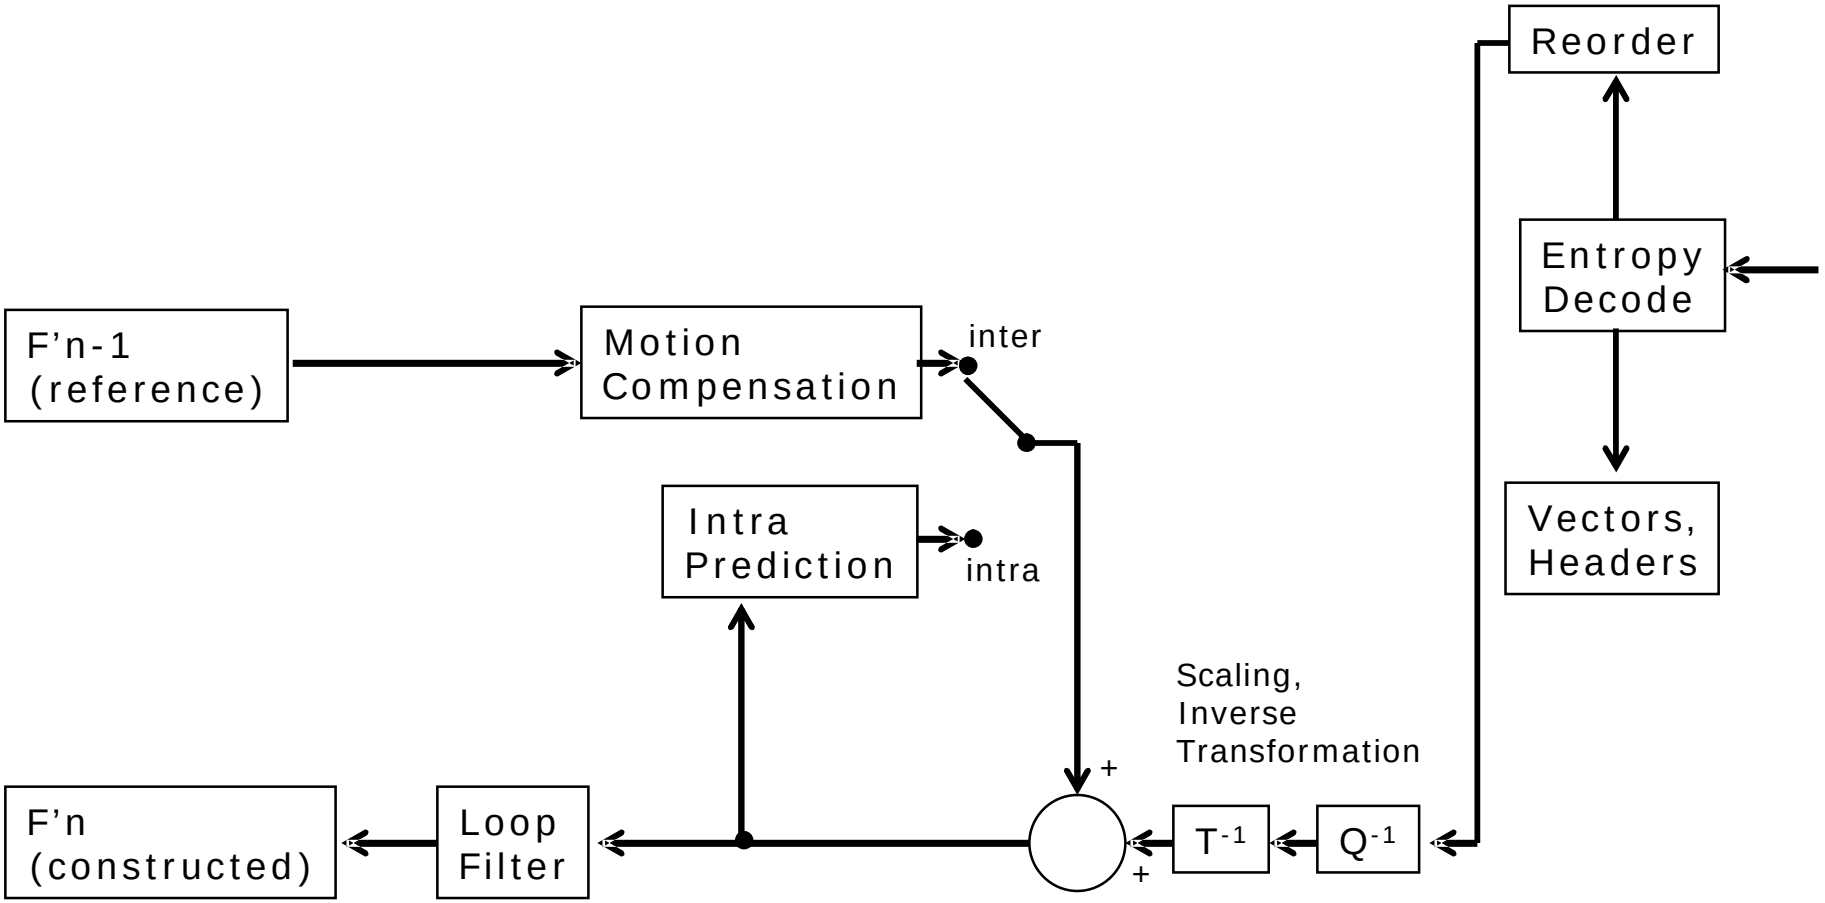

Driving Area Efficiency

- Leverage Decoder ó Encoder functionality overlap
- Programmable HWAs for similar compute functions but with different parameters (such as number of taps), and/or different coefficients
  - DCT/IDCT, 8X8 vs 4x4
  - Quantization, scaling
  - Variable length coding
  - Interpolation (half pixel, quarter pixel)
  - Filtering
- Hardware-Software partition to meet the desired performance and programmability requirements with minimal area

Data driven variability in Decoding

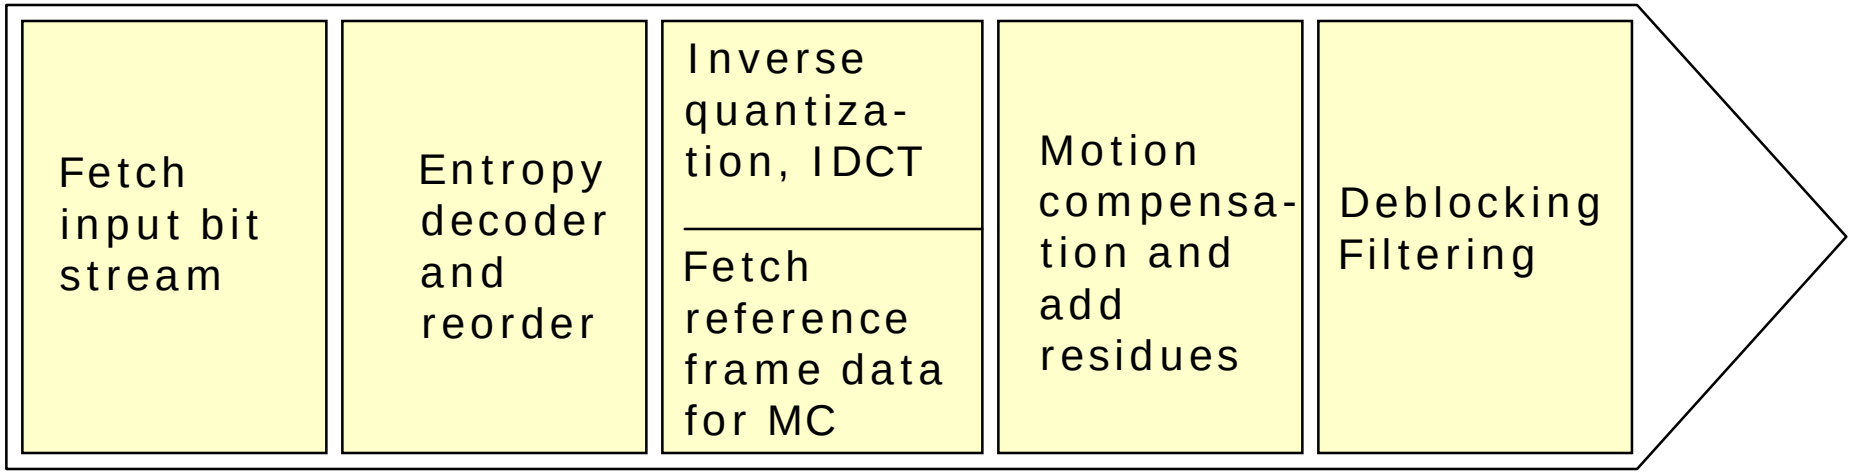

Low Power Design – across all levels

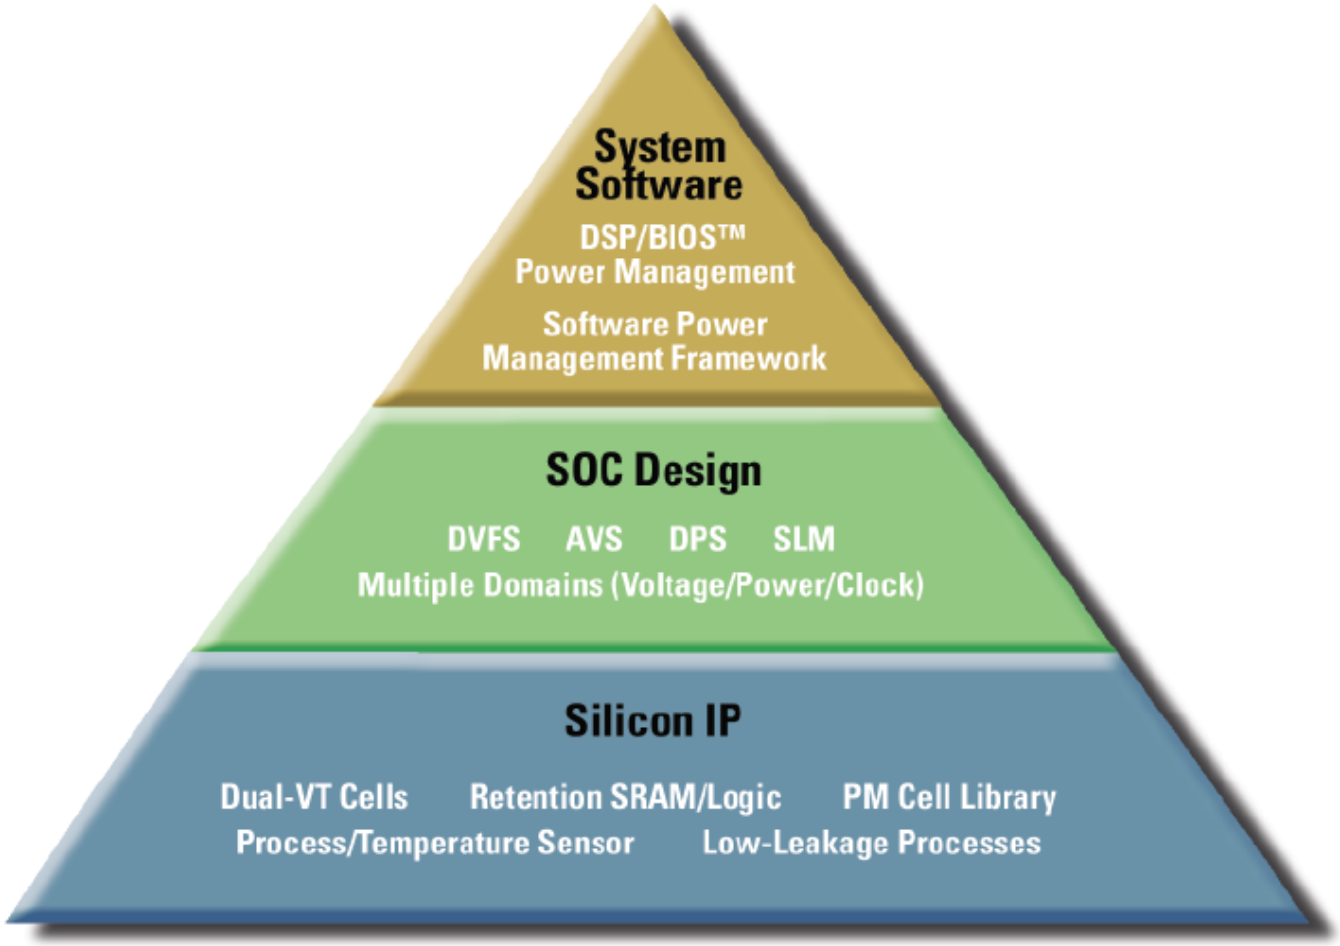

Component level support

Silicon IP

Technology	Description
Retention SRAM and logic	SRAM and logic retention cells support dynamic power switching without state loss, lowering voltage and reducing leakage.
Dual-threshold voltages	Higher threshold for lower leakage and lower threshold for higher performance.
Power management cell library	Switching, isolation and level shifters support multiple domains in SOC implementations.
Process and temperature sensor	Adapts voltage dynamically in response to silicon processes and temperature variations.
Design flow support	Complete, nonintrusive support for easily integrating SmartReflex technologies.

SoC level Power Management Strategies

Technology	SOC architectural and design technologies Description
Adaptive Voltage Scaling (AVS)	Maintains high performance while minimizing voltage based on silicon process and temperature.
Dynamic Power Switching (DPS)	Dynamically switches between power modes based on system activity to reduce leakage power.
Dynamic Voltage and Frequency Scaling (DVFS)	Dynamically adjusts voltage and frequency to adapt to the performance required.
Multiple Domains (Voltage/Power/Clock)	Enables distinct physical domains for granular power/performance management by software.
Static Leakage Management (SLM)	Maintains lowest static power mode compatible with required system responsiveness to reduce leakage power.

Power optimal MHz-Vcc Operating Point

- Lower Vcc helps both dynamic and leakage power
- If Vcc is lowered while keeping MHz same – can result in area increase – impacting cost and negating any power gain
- At architecture level – MHz/Vcc for a given technology drives the degree of parallelism and pipelining
- The choice of target format for power optimization impacts area efficiency – for example, an implementation which gives lowest power for 720P resolution is likely to be different (higher area) than the implementation which gives lowest power for D1

Power Reduction - at Application / Video stream level

- If it's decode function – turn off (clock gate/power down) encode functionality (e.g. Motion Estimation)
- For the standard and the profile to be processed, turn off hardware supporting all other standards and profiles (e.g. if MPEG4, turn off CABAC engine in entropy decoder)
- Dynamic frequency and voltage scaling – set the DV engine frequency and voltage operating points – depending on the resolution being supported – D1 at 30fps requires ~2.66 times lesser compute than 720P at 30fps.

Power Reduction @ Frame level

- Turn off un-used hardware depending on I vs P vs B frame
- Turn off un-used hardware depending on Interlaced vs Progressive content

Power Reduction @ MB Level

- Turn off individual hardware accelerators as soon as the computation for the current Macro-block is done (due to variability, the pipeline cannot be fully balanced)

- During motion-compensation the compute requirements vary depending on 1 motion vector vs 4 motion vectors per macro block, they also vary depending on motion vector resolution in terms of pixel vs half pixel vs quarter pixel.

- Turn off deblocking filter, if boundary strength is 0 or there is significant change (gradient) across block boundary in the original image.

Dynamic vs Leakage Power Scaling with Resolution

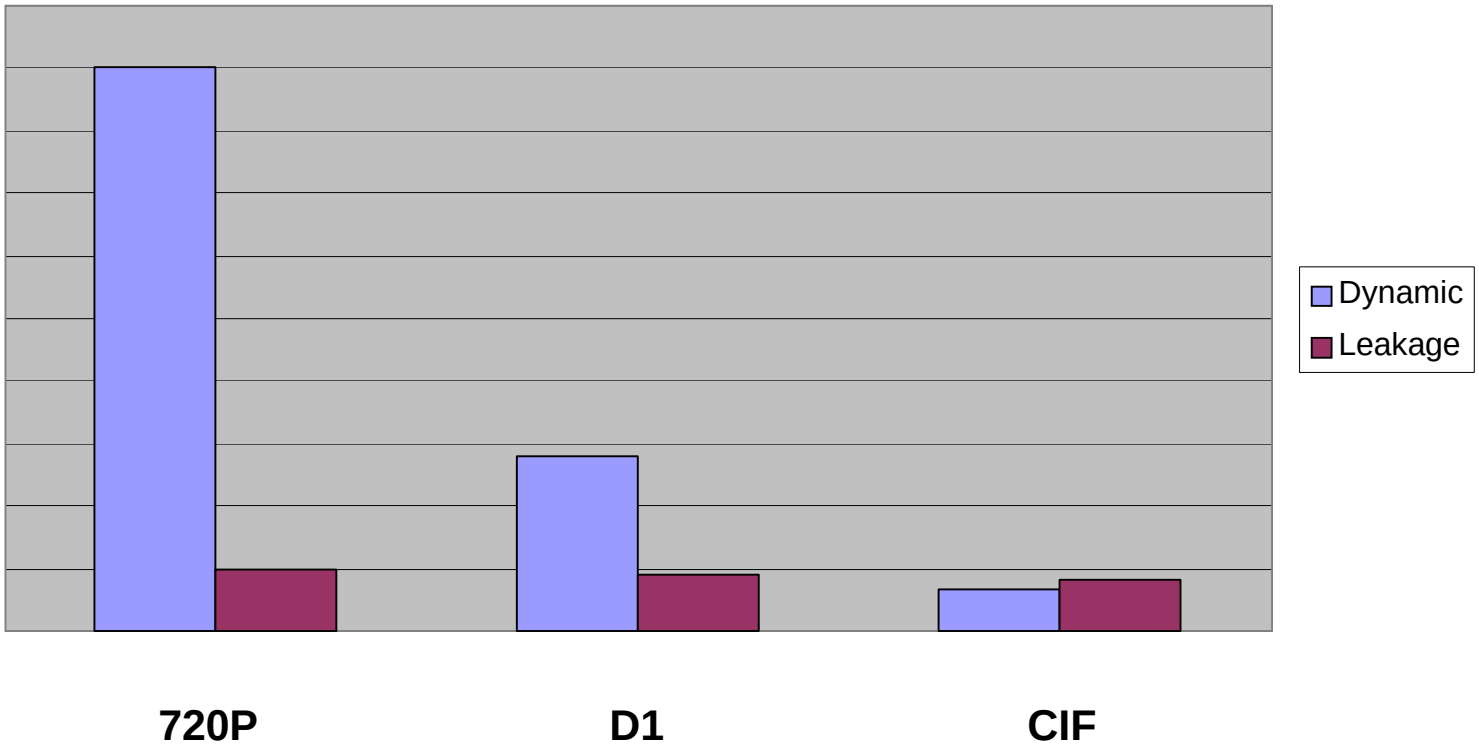

Power Reduction for CIF

- Compute requirements significantly lower, voltage scaling is limited by Vcc-min.
- Running the engine at lower frequency without lowering the voltage – does not help save energy
- Multiple approaches:
  1. Significant cycle overhead in completely switching off the engine and switching it back on – does not help at macro-block level, marginal gain at frame level, but done over a group of frames can give power reduction
  2. Power down the engine but save the state using retention flops and putting memories in the retention mode – area overhead
  3. Design the engine as a “bit slice” and switch off one half while processing CIF – has software implications.

DVFS – applicability at SoC level?

- Audio does not scale with resolution
- Any system function which demands real-time response in a narrow time window
- Modules in the video output processing chain which are tied to the resolution of the display device as against resolution of the video being processed

- Implies multiple voltage domains- can have system level cost implications from PMU standpoint

Managing data bandwidth

- Increasing resolution – implies scaling the IO bandwidth accordingly – but may not be feasible, practical – DDR speed limitations, SDRAM limitation, power, area impact etc.

- Need architecture level solution to address this bottleneck
  - On-chip buffers
  - On the fly computation
  - Improving efficiency of 2D transfers
  - SDRAM data organization
  - Algorithmic solutions?

- At lower resolution, can minimize SDRAM power by powering down unused banks

EDA Challenges and Opportunities

- System level power estimation/modelling
- Power management – synthesis and verification
- Physical design challenges
  - Automated clock gating
  - Physical design aware low power synthesis
  - Multi-Vt optimization
  - Timing closure at multiple corners (with DVFS – need to sweep Vmin and Vmax range)
- Building a configurable IP generator – supporting both run-time as well as compile-time scalability (e.g. building a MPEG4 Decode only engine optimized for power and area, with no software change)

Summary

- Portable video entertainment market needs a multi-format, multi-standard digital video engine with HD capability at low cost and low power
- Highlighted the “variability” in the digital video processing needs including data driven variability
- Discussed the entire spectrum of power management techniques and its applicability to the power minimization of the DV engine
- Highlighted a few system level considerations and their architectural implications
- Finally, presented EDA challenges and opportunity

THANK YOU