Mobile & IPTV Transcoding:
A Total Cost of Ownership Comparison between DSP and x86 based Solutions

August 2008

Authors: Nikos Kyriopoulos, Jeeyoun Lee

Introduction

There are a variety of technical and market forces that have propelled video distribution over IP networks to the forefront of telecommunications industries including telecom, wireless and Internet service providers. The challenge lies not only in the lack of a unified single industry-standard format among content creators, but also in the increasing variety of networked multimedia devices and media delivery technologies. Therefore a reliable, dense and cost-effective transcoding solution is required for content creators and operators to provide TV services to a rapidly growing customer base, without jeopardizing Quality of Service (QoS) and Quality of Experience (QoE).

Currently available video transcoding offerings range from open-source software-based solutions to dedicated transcoding appliances and utilize different hardware architectures such as General Purpose x86-based Processors (GPPs), Digital Signal Processors (DSPs) and Application-Specific Integrated Circuits (ASICs). Historically in the video industry (e.g. in set-top boxes, DVD players, encoders and media servers) GPP-based solutions tend to be introduced first, and as density, speed, quality and cost pressures increase, DSP-based solutions take over until the product reaches its maturity and can be turned into an ASIC.

However merely comparing the processing capabilities of different hardware architectures does not reflect the true use case of a transcoding solution from the deployment stand point and its business impact. Thus, this paper attempts to compare the different offerings from the TV-services operators' stand point by taking into account the Total Cost of Ownership (TCO) for different head-end deployment scenarios.

Two representative and established products have been chosen for this analysis:

- Media Excel's HERA 4000 encoder/transcoder product, based on TI DaVinci DSPs
- 3rd party encoder/transcoder product, based on Intel quad-core GPPs

Finally, apart from the quantitative TCO analysis, this paper also examines a number of qualitative arguments and investigates the implications of the technological advancements in DSPs and GPPs in the video transcoding market.

Total cost of ownership for a transcoding farm

The TCO analysis examines 3 typical scenarios (IPTV, 3G Mobile TV, Broadcast Mobile TV) and takes into account both the cost associated with the acquisition of the required transcoding equipment with software and the operating cost throughout the lifetime of the deployment, including support & maintenance, power consumption and rack space costs.

Both products, HERA 4000 (DSP-based) and the 3rd party product (GPP-based), are configured in order to provide the same or very similar results in terms of:

- Number of input channels
- Number of output profiles/streams per channel
- Video and audio quality
- Feature set (e.g. pre-processing tool set, bit rate constrains, etc)

For the analysis, the following data were taken into account:

- Acquisition cost: retail price

- Support/maintenance cost: 20% annually
- Rack space: co-location rates
- Power consumption costs: consumer rate per KWH
- All monetary figures are expressed in US dollars

IPTV (Standard Definition)

A typical IPTV deployment has the following characteristics:

Table 1: IPTV Scenario

# of channels	Lifetime	Input/Source	Output Profile
120 channels	5 years	MPEG-2 SD (Sports/Movies) MP@ML 4-6Mbps	H.264 SD MP@L3 1.2-2Mbps

Note: High definition channels are excluded from this analysis. (Even though HERA 4000 is capable of accommodating 2 HD channels of H.264 output, at the time this analysis was authored no GPP-based products exist that can deliver HD H.264 output with comparable video quality.)

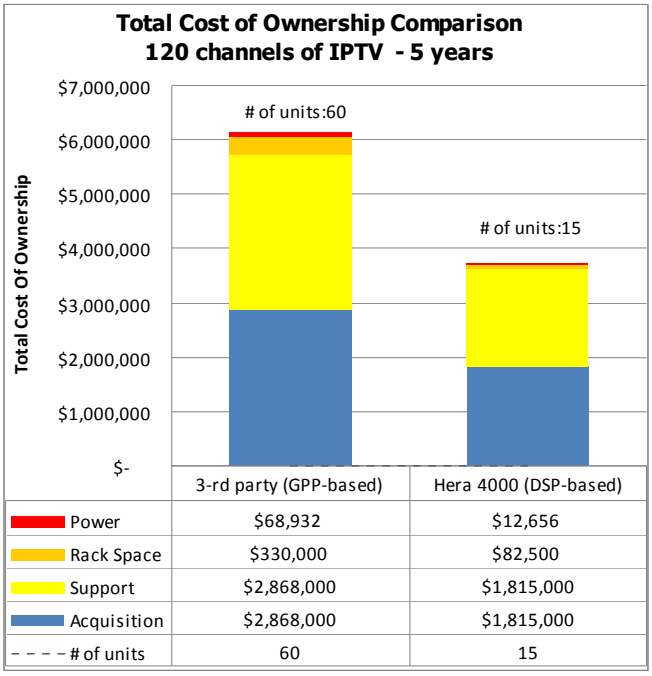

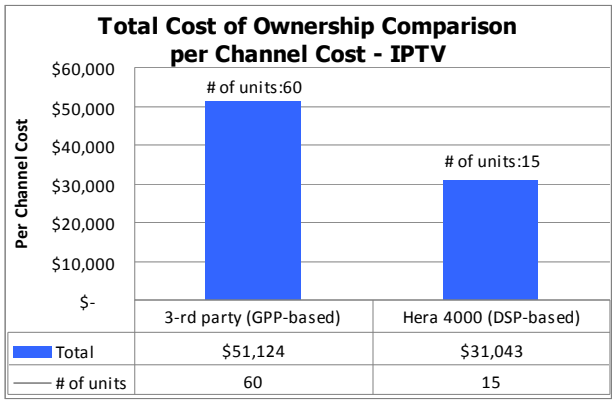

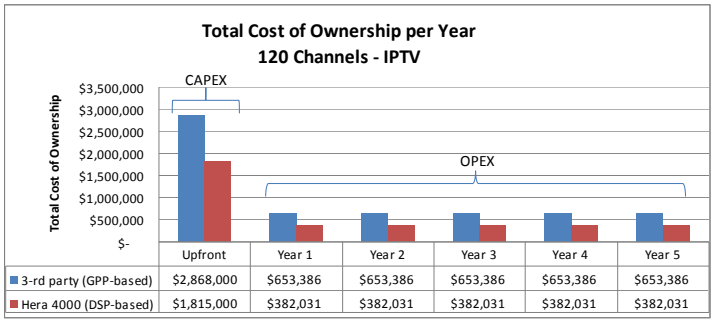

Findings:

- 4 times more GPP-based units are needed compared to DSP-based units (60 vs. 15)
- 58% higher acquisition cost for the GPP-based solution
- 4 times higher rack space cost for the GPP-based solution
- 445% higher power consumption cost for the GPP-based solution
- 1.7 times higher operating cost (OPEX) for the GPP-based solution
- 71% higher total cost of ownership (CAPEX and OPEX) for the GPP-based solution

Mobile TV 3G

A typical 3GPP Mobile TV deployment has the following characteristics:

Table 2: 3G Mobile TV Scenario

# of channels	Lifetime	Input/Source	Output Profiles
48 channels	3 years	MPEG-2 SD (Sports/Movies) MP@ML 3-6Mbps	2x H.264 CIF/QVGA 2x H.263/MPEG4 CIF/QVGA 2x H.263/MPEG4 QCIF

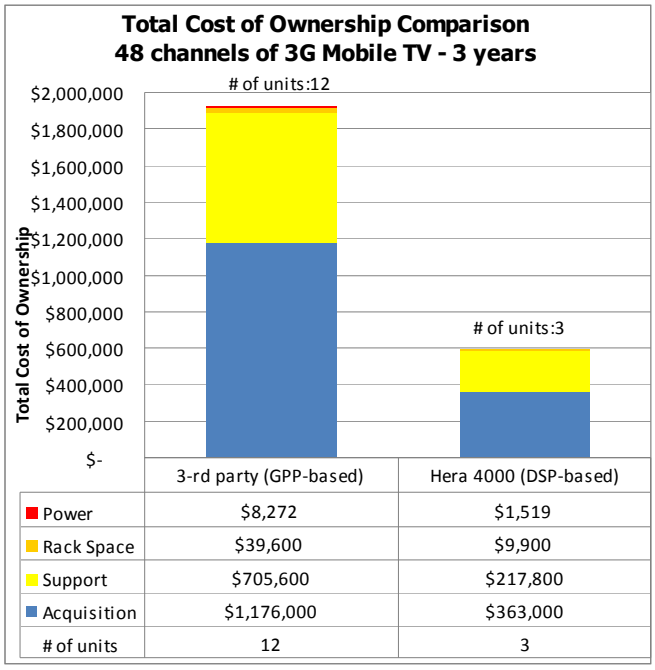

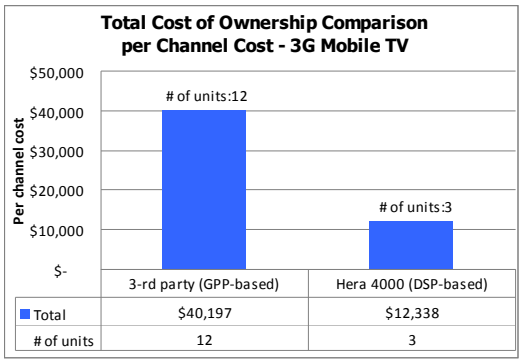

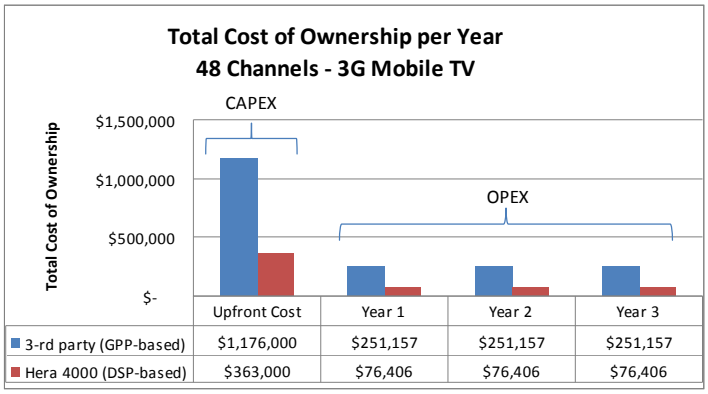

Findings:

- 4 times more GPP-based units are needed compared to DSP-based units (12 vs. 3)
- More than 3 times higher acquisition cost for the GPP-based solution
- 4 times higher rack space cost for the GPP-based solution
- 3.2 times higher operating cost (OPEX) for the GPP-based solution
- 226% higher total cost of ownership (CAPEX and OPEX) for the GPP-based solution

Mobile TV DVB-H

A typical DVB-H Mobile TV deployment has the following characteristics:

Table 3: DVB-H Mobile TV Scenario

# of channels	Lifetime	Input/Source	Output Profiles
16 channels	3 years	MPEG-2 SD (Sports/Movies) MP@ML 3-6Mbps	2x H.264 CIF/QVGA CAS / Stat Mux

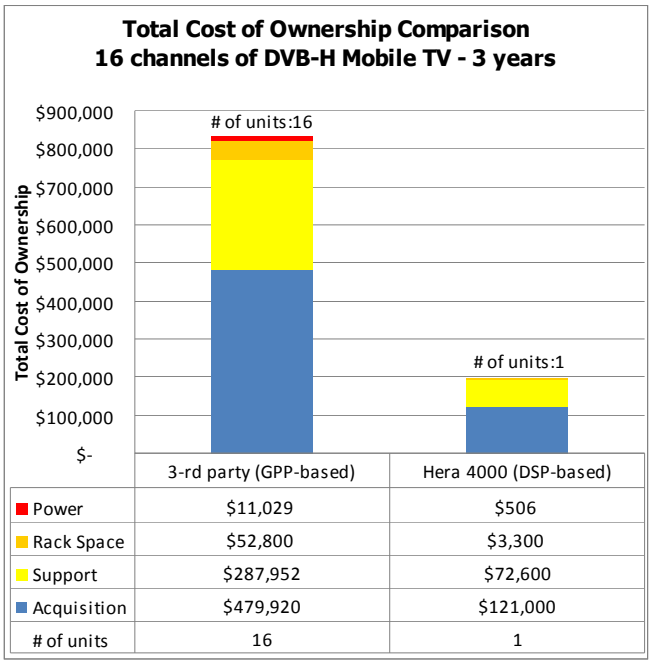

	3-rd party (GPP-based)	Hera 4000 (DSP-based)
Power	$11,029	$506
Rack Space	$52,800	$3,300
Support	$287,952	$72,600
Acquisition	$479,920	$121,000
# of units	16	1

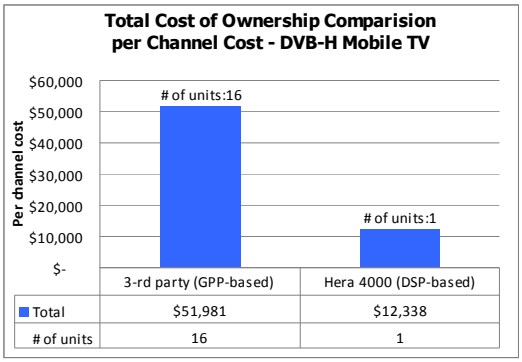

	3-rd party (GPP-based)	Hera 4000 (DSP-based)
Total	$51,981	$12,338
# of units	16	1

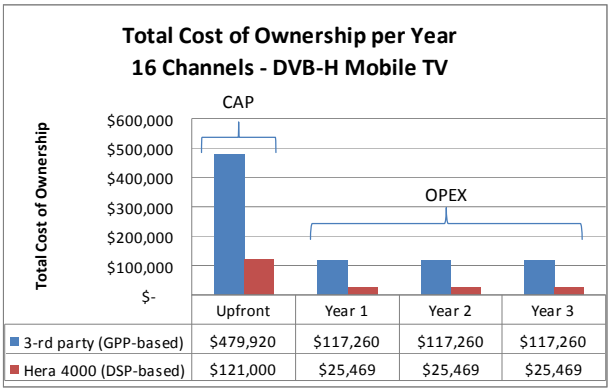

	Upfront	Year 1	Year 2	Year 3
3-rd party (GPP-based)	$479,920	$117,260	$117,260	$117,260
Hera 4000 (DSP-based)	$121,000	$25,469	$25,469	$25,469

Findings:

- 16 times more GPP-based units are needed compared to DSP-based units (16 vs. 1)
- 3 times higher acquisition cost for the GPP-based solution
- 15 times higher rack space cost for the GPP-based solution
- 4.6 times higher operating cost (OPEX) for the GPP-based solution
- 321% higher total cost of ownership (CAPEX and OPEX) for the GPP-based solution

Notes:
Mobile TV service providers tend to replace equipment more frequently than IPTV providers (3 vs. 5 years) as technological advancements provide significant bandwidth and density savings. This explains the difference in OPEX figures between the different scenarios.

Rack space costs are measured based on co-location rates. This does not always realistically reflect the true cost of a deployment, since the cost of expanding a data center to accommodate 60 vs. 15 units (in the case of IPTV for example) is exponentially different.

The TCO analysis does not take into account redundancy and management equipment. This is done for practical purposes only and does not affect the comparison. HERA 4000 provides modular/scalable density and can accommodate all high-availability strategies through a fully-customizable management appliance.

Further Analysis

The TCO analysis above suggests that GPP-based transcoding solutions are simply not competitive in terms of delivering the same number of channels for a given scenario in the same space and with the same electrical power consumption as DSP-based solutions. Apart from the quantitative TCO analysis however, a number of qualitative points need to be taken into account as well.

Quality of Service/Experience

For a video transcoding solution, density and acquisition/operating costs are far less

[1] Tim Siglin, Commentary: Cisco, TI "Go Video", Streaming Media, December 5, 2006

important than Video Quality and overall Quality of Service/Experience. Furthermore, channel density and video quality appear to be two opposing forces as both fight over the processing power available on the system. Often, GPP-based transcoding solutions provide a minimal set of pre-processing or encoding tools in favor of higher channel density. On HERA 4000 however, DaVinci DSPs lift this trade off and allow for significantly higher video quality by enabling features such as motion adaptive inverse telecine, de-interlacing, noise reduction and "real-time" multipass encoding.

Time-To-Market – Versatility

General Purpose Processors by definition are suited for general purpose use and thus the development cycle of general purpose software tends to be easier and shorter. DSPs, on the other hand, are focused towards time-critical media processing and real-time applications which make them ideal for video transcoding. Development cycle on DSPs tend to be longer mainly due to limited availability of engineering talent, however, Media Excel's team of codec engineers is highly qualified to take advantage of DaVinci's feature set and deliver highly optimized transcoding solutions. Furthermore, DSPs bridge the gap between the versatility of GPPs and the specialization/robustness of ASICs, allowing for solutions that are both field-upgradable and also dense and robust. This shortens time-to-market and enhances products' life span.

Scalability

As more and more media-aware consumer devices emerge, transcoding needs become more and more prominent for the end-user. Even though currently media transcoding occurs mainly on the head-end or at the regional level (e.g. Content Distribution Networks - CDNs), major industry players[1] are investigating the advantages of edge transcoding. Transcoding at the edge enhances user's experience by providing better customization and targeting (e.g. device-specific content repurposing, targeted advertisement, content mobility), significantly reduced bandwidth needs at the head-end, but could also increase the infrastructure costs. In this scenario, a highly-dense and scalable transcoding solution is required. DSPs present an inherent efficiency edge of 2x in terms of footprint and 8x to 15x in terms of power consumption compared to x86 CPUs. Furthermore, heat dissipation and peripheral modules (SBC, memory, etc) inflate exponentially the real estate occupied by a GPP-based media processing solution, effectively disqualifying it from the edge-transcoding race.

Hardware Upgradability

Advancements in processing power dictate frequent hardware upgrades, however in a DSP-based solution this translates to the exchange of the DSP modules only, while in a GPP-based one it requires the exchange of the entire unit. Taking this into account, a Mobile TV deployment upgrade based on HERA 4000 will cost a quarter (~27%) of that of a GPP-based solution (assuming no change in prices). A blade-based solution makes good sense for high-capacity easily-upgradable deployments, however the footprint and heat dissipation needs of a GPP-based blade still can not compete with a DSP-based one.

Technology Roadmap

TV service providers are effectively shareholders of the technology they decide to invest on, as they require assurance about the viability, upgradability and generally the future of their sizable investment. It is important therefore to examine the level of commitment each technology vendor is giving to media processing. TI for example exhibits by far the largest patent portfolio in media processing compared to those of Intel's or AMD's. TI's focus and core business has always been towards facilitating signal/media processing and this provides substantial reassurance for the company's future commitment. On the other hand,

General Purpose Processor vendors by definition have much wider area of interest and their commitment to supporting media processing varies according to market and competitive forces.

Outlook Even though the analysis above suggests that DSPs are best suited for media transcoding solutions, it is important to ensure that Moore's law won't shift this balance. The graph below summarizes the evolution in processing power for both GPPs and DSP clusters.

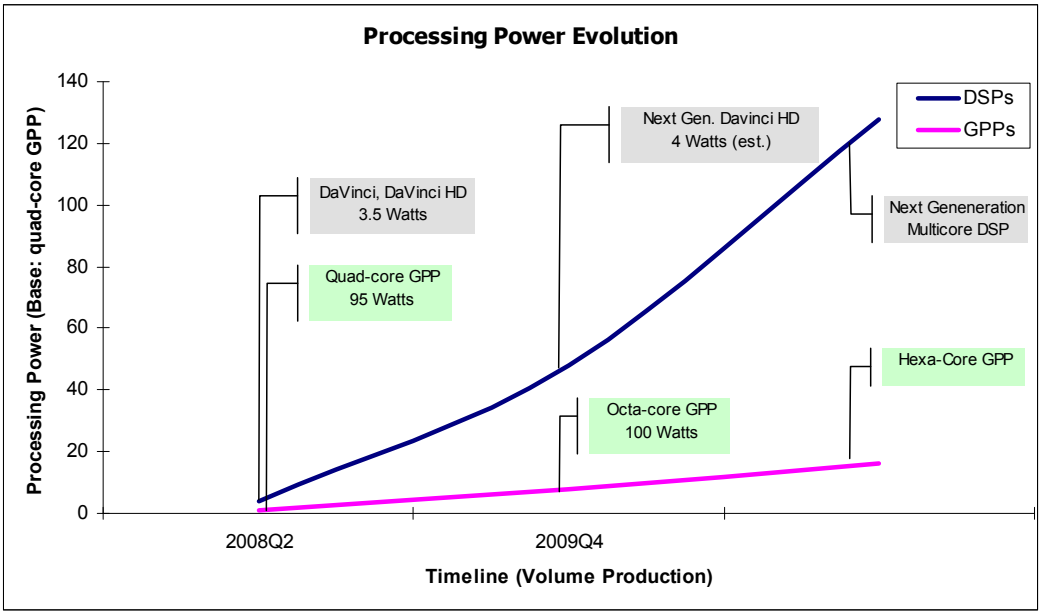

Conclusion The analysis above concludes that, for the foreseeable future, DSPs remain far more suitable for media transcoding solutions than GPPs, both in terms of deployment cost but also in terms of technological investment.