Home media center built with Linux on PowerPC

By Sergei Larin
Applications Engineer
Freescale Semiconductor Inc.
E-mail: sergei.larin @freescale.com

Pieter Van den Abeele
Director of Software Applications
Genesi USA Inc.
E-mail: pvdabeel @genesi-usa.com

A major aspect of the market competition for data flow control to future homes is the consolidation of various functions–TV, phone, Internet, media content–into a single multimedia center that can handle all these seamlessly.

To ease and standardize the flow of data into the home, the delivery method and management of services are crucial considerations. In many cases, delivery depends on two-way communication and guaranteed control of proprietary media content.

This article focuses on handling the multimedia content after it is delivered. The various data streams to consolidate and manage range from voice and VoIP, cable/satellite TV with DVR functionality, and Internet services to on-demand media delivery and online gaming.

The wide variety of formats and the potential need to interchange them likewise pose unique challenges for center or a multimedia home platform. In this article, the term multimedia center designates the actual hardware (center), and the multimedia system integrates the software component of the design.

Traditionally, each unique media function is handled by a dedicated device and a dedicated communication line. The result is the duplication of, for example, a digital cable for TV programming, an Internet provider with IP phone service and a regular phone line–all using essentially the same infrastructure, but different delivery methods and dedicated devices. The trend is to consolidate multiple functionalities into a single device. A home computer often plays this role, but in a consumer-friendly way, the challenges of price, simplicity and efficiency remain.

To address these challenges, this article presents a multimedia center implementation using software techniques. It considers PowerPC architecture as a candidate and provides performance measurements to support these claims.

Current implementations

As an example of the consolidation trend, let us point to several recent offerings in the area. Several solutions built around the Intel Viiv digital rights management (DRM) platform have been announced recently–PC Fusion 2, Ace LH series of Digital Entertainment systems, AVidius Media Systems. All of these closely resemble a general computational platform with a software layer to provide multimedia handling capabilities and a unified interface. The general purpose dual-core CPU (Core Duo) is used in this model.

Another example is Viiv's direct rival, AMD Live!, with a very similar approach. Both platforms illustrate an implementation based mainly on a software solution with substantial computation power dedicated from a general-purpose core. At the same time, both platforms incorporate DRM and will be using Microsoft Windows OS (for instance, Windows XP Media Center Edition).

Earlier system examples include the following:

- Thompson IP900, introduced back in 2003, is a MIPS-based SoC, while IP1000 is based on Intel x86. Both systems are nominally IP STBs, but the IP1000 decodes MPEG-4 and H.264, and the IP900 handles only the MPEG-2 stream (**Table 1**).
- Samsung SMT-7000S is based on Intel Celeron and is marketed by Samsung as an IP/DSL STB.
- Fujitsu Siemens Series 200, 300 IP-STB uses a 1.25GHz Intel CPU.
- Handan uses the Celeron 733MHz or 1GHz CPU.
- Humax uses the VIA Eden CPU (800MHz).
- Softier IP STB based on TI DSP and Linux OS.

All of these examples illustrate an approach that uses a general-purpose CPU with limited peripherals to implement a multimedia center. On the other end of the spectrum are numerous systems with an extremely light-weight general-purpose core to execute control functions and with several peripherals serving as a dedicated hardware solution. These systems represent today's vision of a collection of many dedicated components, each executing a specific function. These platforms target the high-volume, low-price market and tend to be very rigid in their implementation.

Feature set

There is a wide design space for multimedia systems and general disagreement about their definition and classification.

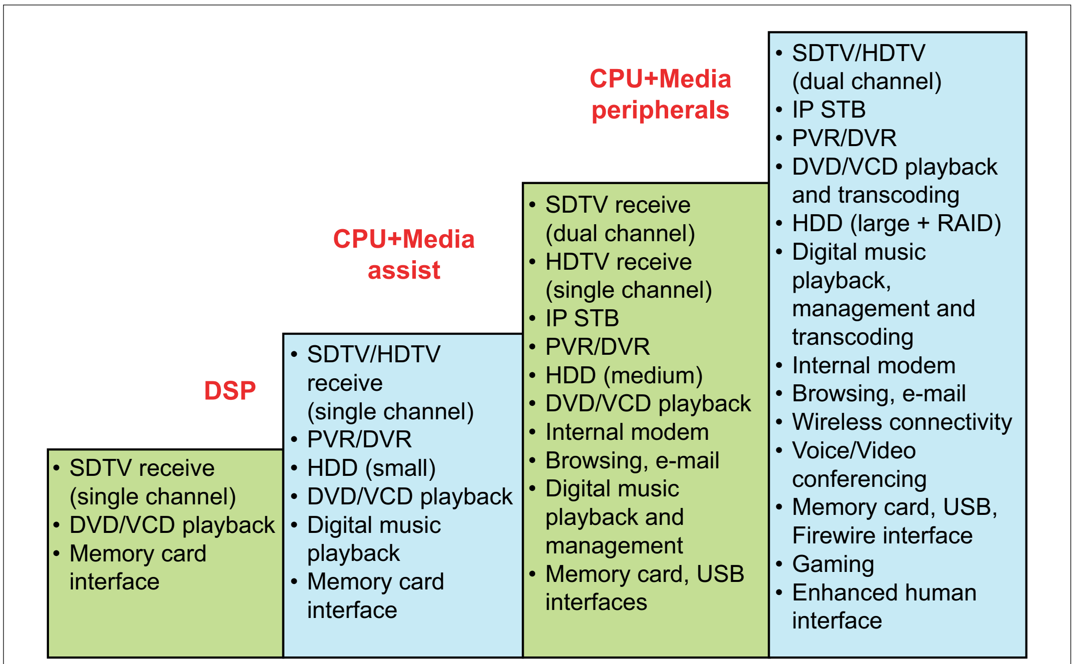

Figure 1: The design space of a multimedia system can be loosely classified into several tiers according to implementation, cost and feature set.

The design space can be loosely classified into several tiers according to implementation, cost and feature set (**Figure 1**). DRM is a part of any system by default.

To set up a good tier 3 multimedia system, the desired feature set as follows:

- STB with digital video recording functionality–It should be able to receive (terrestrial, satellite, IP TV carriers), view, pause, rewind and record live TV (timeshift). The system handles two SDTV streams simultaneously, one viewed (de-interlacing, if needed) and another recorded or overlaid as a picture-in-picture. There is programming guide retrieval with advanced digital video recording management: schedule recordings ahead of time, following programming changes and resolving conflicts between multiple recordings.
- DVD, CD, VCD playback, management (transcoding) and writing/burning
- Media cataloging and reliable storage that can seamlessly interface to NAS or implement an on-board RAID system–Redundancy-enabled storage such as RAID can be a significant marketing point, but is also very sensitive to the form-factor of the system. For the current design, simple high-capacity HD is used for content storage. RAID is targeted as future work.
- Internet access console that includes browsing, automated retrieval of information, such as weather and news headlines, and elements of data management, including headline analysis for keywords of interest and offensive content filtering–Functionality may include networking appliance capabilities: in-house router functionality, software firewall maintenance and virtual private network (VPN) support. The proposed system handles only automated data retrieval, which depends on external content providers.
- On-demand delivery of media, such as movies, music and games–The delivery requires DRM and external provider interaction.
- Gaming support in both the on-demand and game console modes. This option may demand major changes in system design, thus defining many other aspects of the system, such as the level of graphic card support. We use the system only as a game console and implement several open source games.
- Voice- and video-over-IP functionality, such as video conferencing or simply digital phone service using an IP connection–This functionality puts a specific demand on a variety of protocol support, both for connection management and for voice/video codec (SiP, H263, G711, G729 etc). This feature is fully implemented and allows conducting a full duplex video call.
- Primitive artificial intelligence functionality–This is a bonus to improve human interface and feel. It can assist in searching a program guide based on close caption or programming guide descriptions. It enables easy, hands-free voice recognition and text-to-speech, and navigation among features. The text-to-speech feature is implemented, but not currently integrated with the rest of the system. The artificial intelligence component is also reserved as future work.

Most of these features require high computational power and call for a powerful CPU and software solution. The form factor for such a system, in addition to likely in-room positioning, results in a noiseless (fan-less) low-power setup. Given these requirements, the ratio of processing power and power dissipation to cost is very important.

Although many of these functions overlap with the role of PCs today, in a multimedia center, they are limited mainly to consumption–not creation–of the media, non-modular design, and emphasis on a seamless human interface.

Open-source approach

The open-source approach widely facilitated by Linux OS is a step away from high-budget implementations and proprietary operating systems, and a step towards cooperation and collaboration. The open-source community provides access to implementation of various algorithms and standards, enabling an excellent starting point for a design.

Various open-source projects implement multimedia server functionality. Most are collections of multiple features, but two deserve inspection–MythTV and Freevo. Both projects implement a digital video recorder for an STB-like design, and have certain advantages and disadvantages.

After both applications were tested in independent implementations, MythTV's design emerged as more configurable. Additional plug-ins/components can be integrated seamlessly into the common GUI front-end, which in turn, is easily customized to meet vendor requirements.

MythTV has the benefit not only of design flexibility, but also easy distribution between the main server and multiple satellite stations. The main server runs MythBackend with MySql database and other high-intensity jobs, while a satellite station runs the MythFrontend thin client, which receives only processed content and outputs it to a user console (regular TV or progressive monitor).

Thus, the actual compute-intensive processing can be physically partitioned from the access console, which is important when too much processing power and features are concentrated in a single location. In this scenario, a backbone system with RAID storage and a powerful processing core can reside outside the entertainment room, with light profile satellite stations located throughout the house.

Open-source multimedia server projects are sometimes criticized for their installation complexity and high initial investment in the software configuration. These arguments are voiced by Linux enthusiasts who configure the software for their custom systems. In a production environment, these criticisms don't apply. However, updates and maintenance of previously released software components require attention to the version of the underlying Linux OS.

Multiple vendors pack free Linux code and redistribute it under various brand names and different packaging options. Such packages include preselected (and often precompiled) sets of features/applications.

One notable exception is Gentoo Linux, which markets not a distribution but a system knowledge tool that allows designers to create their own custom distributions tailored

CPU	Thompson IP900 MIPS	Thompson IP1000/1001 Celeron
Clock rate	242MHz	733MHz
Memory	DDRAM 64MB	DDRAM 64MB
OS	VxWorks	Linux, WinCE
Video Codecs	MPEG2	MPEG-2, MPEG-4 ASP, WM9, H.264

Table 1: Thompson IP900 and IP1000 systems are nominally IP STBs.

to a specific platform and need. The designer is in charge of selecting the required components, and the system knowledge tool guarantees that all the appropriate packages are delivered and all dependencies are respected.

The high-level end application (such as MythTV) check out prompts an automated yet very complex process of verifying components and package presence on the system, their interdependency, necessary source code delivery, and a custom build followed by installation of required components—all from a single command.

Development time is minimized. Moreover, there is a guarantee that any future updates and/or bug fixes in any of the components prompts updating and rebuilding as needed. The result is an improved version of the final product. With a multimedia system built with a general-purpose CPU at a known location, adding new features and fixing problems become very similar to the software update mechanism used by most commercial software vendors.

Platform-specific

In implementing the multimedia server, the main advantage of PowerPC G4/G4plus processors is their excellent performance-to-power ratio. The PowerPC architecture can use passive cooling to provide zero-noise, fan-less implementation that enables high processing power for multiple HDTV decoding streams in software, or a video conference while recording an HDTV stream.

The PowerPC architecture is inherently friendly to multimedia content. All G4/G4plus processors offer the advantage of AltiVec single instruction-stream, multiple data-stream (SIMD) unit. This AltiVec technology was originally developed by a design consortium consisting of Apple Inc., Motorola SIMD SPS and IBM Inc. to accelerate multimedia processing.

The AltiVec technology augments the processing power of a sophisticated superscalar core by offering additional vector/parallel programming capabilities that can significantly speed up many classes of communications and multimedia applications: IP telephone gateways, multichannel modems, VPN systems, speech processing systems, echo cancellers, digital image and video processing systems, scientific array processing systems etc.

The 128bit AltiVec engine remains the most powerful option for SIMD extension processing to a general-purpose superscalar core. Numerous applications are written for it, and most trade-offs and design nuances are well-understood. All new G4/G4plus cores from Freescale include AltiVec technology, which is fully backward-compatible with all the software developed for it.

To illustrate the potential of AltiVec, we refer to the Embedded Microprocessor Benchmark Consortium (EEMBC), which has published multiple benchmarking scores in many application categories over the years. Of interest is the Digital Entertainment subcommittee, which contains standardized benchmarks for MPEG-2 and MPEG-4 encoding end decoding, MP3 decoding, JPEG processing etc.

This subcommittee has produced a total of 15 benchmarks with 69 data sets grouped for convenience in several so-called sub-marks: MPEG Decode Mark, MPEG Encode Mark, Encryption Mark and Static Image Mark. There is a single aggregate score for the entire 69 benchmark/data set permutations, known as Digital Entertainment Mark (DENMark).

The MPEG_Decode_Mark is composed of scores for MPEG-2, MPEG-4 and MP3 decoding. The MPEG Encode mark is composed of MPEG-2 and MPEG-4 encoding scores. Static Image-Mark is a collection of Compress/Decompress JPEG image and several color processing filters. Freescale processors have two score variations:

- Regular—It is known in EEMBC as out-of-the-box (OOB) and used by default for all certifications. It is obtained by a simple compilation of the standard C code with a fixed set of compiler options.
- Optimized (OPT)—It adds core-specific intrinsics to the C code.

The difference between OOB and OPT is what we want to concentrate on—it shows the performance advantage of AltiVec technology and yields a 217 percent average performance increase for the entire benchmarking suite.

For the Freescale products, typical power dissipation while running this code measures at about 20W for MPC7447a at 1.4GHz, and about 15W (L Spec) for MPC7448 at 1.7GHz. This low power dissipation is well within the limits for a passively-cooled simple heat sink solution. The AltiVec technology measures at a negligible five percent of the total power dissipation while delivering a several times performance boost.

This amalgamation of a high-performance superscalar CPU with the 128bit wide parallel SIMD engine, along with low power dissipation, makes the general-purpose PowerPC core very attractive for multimedia server implementations.

Most algorithms in our system use AltiVec technology in different degrees, which enables the design and permits the use of a much slower CPU for the function.

SIMD for multimedia

One of the most demanding algorithms to be encountered by the multimedia system is software MPEG decoding for broadcast TV and DVD playback. The performance numbers from EEMBC MPEG-2 represent a DVD playback in both resolution and decoder configuration. MPEG decoding for TV broadcasting is very similar.

Generally, MPEG-2 is most commonly used as the compression algorithm for both SDTV and HDTV broadcasts. Although MPEG-2 supports up to 4:2:2 YUV chroma subsampling and 10bit quantization, which EEMBC uses in most data sets, DTV broadcasts use only 4:2:0 and 8bit quantization to save bandwidth, which in turn, changes CPU workload slightly. Some HDTV broadcasters also plan to use MPEG-4, which offers a slightly higher compression rate for a given perceived quality and/or possibly multiple channels to deliver a high-end HDTV signal.

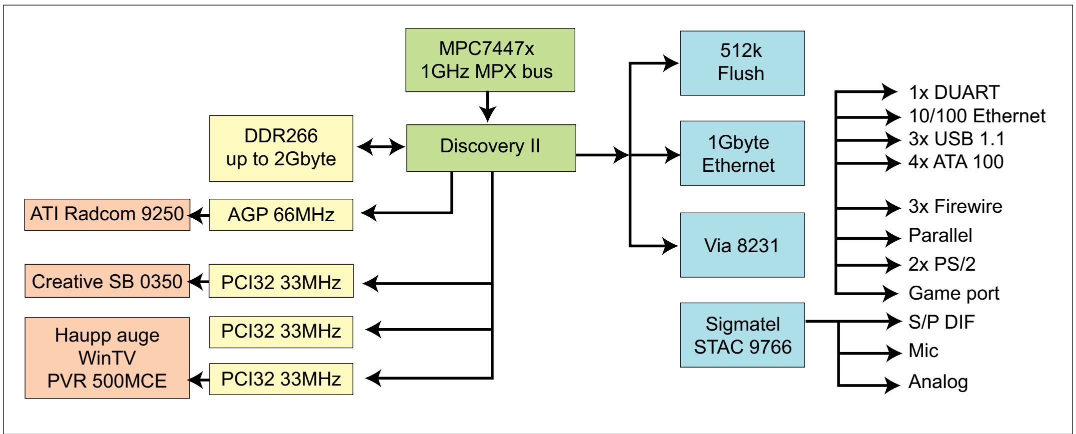

Figure 2: All design and implementation details (gerber files, schematics, component list, and so on) are freely available, and with the open source software layer, add up to a truly open design.

Format index	Type	Vertical resolution	Horizontal resolution	Aspect ratio	Scan type	Refresh rate (Hz)	Raw data rate (Mbps)	MPEG-2 MP ML Data rate (Mbps)
1	SDTV				interlaced	30	105	2
2		2				4	169	3
3		3				0	211	3
4			640	4:03	progressive	60	422	6
5					interlaced	30	116	2
6		2				4	186	3
7		3				0	232	3
8			704	4:03	progressive	60	464	7
9					interlaced	30	116	2
10						24	186	3
11						30	232	3
12	EDTV	480	704	16:09	progressive	60	464	10
13						24	506	8
14						30	633	9
15		720	1280	16:09	progressive	60	1,266	16
16					interlaced	30	712	11
17						24	1,139	17
18	HDTV	1080	1920	16:09	progressive	30	1,424	18

Table 2: The North American Advanced Television Standard Committee defined 18 formats for DTV.

The North American Advanced Television Standard Committee (ATSC) defined 18 formats for DTV (**Table 2**). Here, the MPEG-2 data rate is a close approximation and can vary slightly. Also, notice that the actual NTSC signal has some overhead for audio and transport organization so the end data rate varies to some extent. HDTV is using the Dolby Digital (AC-3) format to support 5.1 surround sound, theater quality audio track. Nevertheless, the data rate for audio track is negligible compared to the rate of the video signal.

MPEG-2's Main Profile, Main Level decoding in software can be roughly separated into four operations, some of which are locally data-independent and can be executed in parallel (**Figure 2**). As the EEMBC scores in **Table 3** show, AltiVec technology can be used to perform these decoding operations to accelerate individual algorithms. For example, AltiVec unit can accelerate 2D IDCT in a stand-alone mode by nearly 500 percent.

Dequantization (iquant) and motion estimation can each be accelerated by about 200 percent. When all components are integrated for TV MPEG-2 decoding, the overall speedup (an average of 10 data sets, different from EEMBC) is about 245 percent.

This increase is in line with the performance gains reported by the EEMBC set of benchmarks for DVD-like MPEG decoding. MPEG-2 encoding is much more work-intensive than decoding and presents more opportunities for vectorization with AltiVec technology. For encoding, the average speedup is measured at around 480 percent. Encoding is a factor in multimedia systems when transcoding is performed between various formats. In a similar analysis for digital audio decompression (playback) tested for an MP3 algorithm with DVD quality, AltiVec unit delivers an average 220 percent speedup compared to the original implementation.

These figures underscore the effectiveness of AltiVec technology in multimedia format processing and the overall suitability of the PowerPC architecture for an integrated home media center.

Design considerations

In a reference design with real-time performance analysis and workload balance trade-offs, the bulk data flow dictates all design trade-offs. A general-purpose CPU with a software-oriented implementation is considered more flexible and application-independent. Benchmark scores justify the use of PowerPC core with AltiVec technology in the design. This system is then implemented to address the cost/feature trade-off and conduct system-level benchmarking.

The system shown in Figure 2 is an off-the-shelf design for a general-purpose Linux desktop PC that is not specifically targeted for multimedia applications, so some trade-offs are not addressed. As a proof of concept, however, this system serves the purpose. The system, designed by Genesi Inc., is the Pegasos Open Desktop Workstation (ODW) platform. This means all design and implementation details (gerber files, schematics, component list) are freely available, and with the open source software layer implementation, add up to a truly open design.

Function	Time (%)
Parsing bitstream	17
IDCT	14
Reconstruction	32
Dithering	25
Misc. arith.	10
Other	2

Table 3: AltiVec technology can be used very effectively to perform these decoding operations.

This design has the functional latitude to configure applications to explore potential design space and highlight important design trade-offs. One example is the use of the CPU for graphical rendering, minimizing the requirement for video card functionalities. Indeed, the multimedia system mainly addresses 2D graphics with limited support for gaming applications and simple 3D graphics (based on OpenGL). Therefore, the simplest video card in a frame buffer mode (just capable of displaying memory buffer content without processing it) can be used.

One of the most important trade-offs in such a system design is where the MPEG-2 TV signal decoding takes place. To guarantee that the TV signal is distributed freely throughout the system built around PCI bus (PCI 32bits at 33MHz bandwidth = 133Mbps = 1Gbps), the raw SDTV whose data rate is about 105Mbps, is delivered by the TV tuner in compressed form.

The TV tuner used in this case is Hauppauge WinTV PVR 500MCE. The compressed data rate is about 2Mbps for SDTV and up to 18Mbps for HDTV (**Table 2**). Of course, the compressed data must be decompressed/decoded prior to rendering, a function frequently outsourced to a co-processing unit. In addition to decompressing, the following steps may be necessary:

1) Deinterlacing to display the data on a progressive screen, a time-consuming algorithm that can consume a non-trivial amount of processing power;
2) Adding overlay information;
3) Blending with another stream for a picture-in-picture. Note that the picture-in-picture feature requires two TV tuners, which are available at the Hauppauge WinTV PVR 500MCE and simultaneous decoding of two MPEG-2 streams.

Another option is to take the MPEG-2 encoder output from the tuner card, use the DMA controller to transfer the data directly to the ATI Radeon 9250 built-in MPEG decoder, and output the decoded video in interlaced form for display.

DMA use minimizes the amount of processing required from the CPU, freeing the CPU to perform other tasks.

The disadvantages of this approach are the interlaced output unless additional hardware is used, difficult overlay and picture-in-picture construction, decoding of only one stream at a time, and the need to use the Radeon 9250 or another video card with an MPEG decoder in the final design. The preferred option is to allow the PowerPC core to do all the work and use the Radeon 9250 in frame buffer mode only, to be potentially replaced by a cheaper video card.

Contrary to the situation with the graphic card, the Creative Sound Blaster Audigy2ZS (SB 0350) audio card is used to its full potential. It delivers a true 5.10 channel surround sound experience for appropriate media. The software sound manager is Alsa.

DVD playback algorithm is similar to TV decoding. The player for our system is Xine, which is built using the FFMPEG project. For CD, VCD and other video file formats (MPEG-4 derivatives), playback does not represent a significant design challenge and can be handled by the same codec.

Transcoding on the other hand, might require non-standard output formats and/or resolutions, but due to the nature of the transcoding process, it does not have to be done in real-time. Hence, critical performance requirements are relaxed in this case.

Another interesting aspect of this design is the seamless integration of video, VoIP and instant messaging into the system. MythPhone is the default plug-in used in MythTV to provide this functionality. Its feature set includes the SIP communication protocol, G.711 or GSM for voice and H.263 for video calls.

Another option is the Gnomeeting project, which is based on the OpenH323 Library. It supports numerous formats, including H261 (video), G711, G726, G723.1, GSM (audio) and a variety of connection protocols. Gnomeeting is also fully interoperable with Microsoft NetMeeting.

In summary, a 1GHz PowerPC core with AltiVec technology is capable of handling two SDTV data streams simultaneously, with CPU cycles left to spare. Future work will implement HDTV functionality on the same or a similar system, with RAID integration and human interface improvements.