

by Kevin Tanaka and Cindy Kao

AUTOMOTIVE FPGAS INCREASE ON-CHIP SYSTEM LEVEL INTEGRATION, LOWER BOM COSTS

By leveraging PLDs, automotive designers can realize immense cost savings through faster time to market, system integration, design flexibility and a scalable production platform. Today's PLDs offer complete qualified solutions designed specifically to address the many challenges encountered in the automotive industry. The latest FPGAs can handle a variety of I/O signaling standards, including LVDS and RSDS, which give the designer the capability to implement TFT, LCD or OLED into the application they are building.

Automotive manufacturers continue to improve in-vehicle comfort, safety, convenience, productivity, and entertainment, which, in turn, drive the use of diverse in-vehicle digital technologies. However, the long development cycles in automotive are difficult to match with the latest technologies, especially with the constantly changing in-vehicle networking specifications, and the

quickly emerging and disappearing technologies from the consumer market, resulting in high engineering costs and obsolescence. Add low cost targets, extended temperature ranges, high reliability and quality, and limited physical board space to the mix, and the challenges in automotive design are daunting at best. Programmable logic devices (PLDs), such as field-programmable gate arrays (FPGAs) and complex PLDs

(CPLDs), have appeared on the scene and are proving to be a flexible, cost effective and viable technical solution while offering better time to market than traditional hardware solutions currently being used.

The commercial aspects of automotive design are becoming increasingly important. In a Harvard University study based on 391 different-sized designs, it was found

that the average ASIC SoC design takes 14 to 24 person-months, while the average FPGA design runs between six to 12 person-months. This is an average difference of 55% in development time, allowing for speed to market for time-critical designs, while also reducing engineering costs and overhead. Another major factor

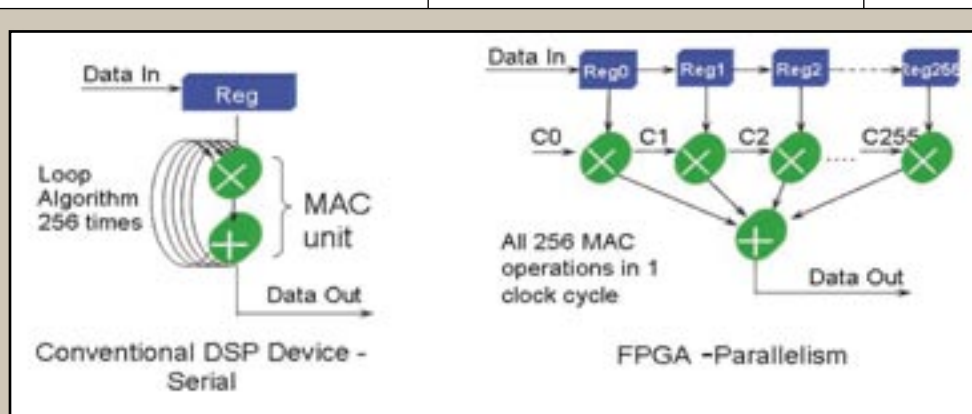


Figure 1. Conventional vs. FPGA digital signal processing.

that is often not calculated into the development cost equation is NRE and mask charges. The average cost for an ASIC SoC mask set ranges between 1 million and 1.5 million U.S. dollars at the 90-nm process technology node, and these costs are doubling with each process shrink. Concurrently, due to the complexity of finer geometries, the chance of having to re-spin an ASIC SoC design due to bugs or layout problems also increase significantly. These two issues in combination must be looked at by the design engineer as a potential risk and added cost. This may be one of the key reasons ASIC design starts worldwide have decreased by about 50% between 2000 and 2003, and continue to decline yearly.

PLDs like FPGAs and CPLDs, on the other hand, provide hardware flexibility due to reprogrammability. Thus, developers are offered the benefit of updating designs from prototype all the way through the production phase. Since PLD designs are programmed using a software bitstream, making quick design modifications are easy and straightforward, and there are no NRE or mask costs.

Since PLDs are scalable in logic density and package migration, they allow designers to make whole-scale changes and still target to the correct pin and logic density. This leads to excellent price per logic cost points and specific tailored pin counts for each design. PLD designs use hardware description language (HDL) to implement logic and C source files for embedded processors. These design source files can be used to target and reconfigure any PLD, any number of times. Designers may also leverage existing designs or take specific parts of designs for re-use in new projects.

This scalability and re-use of code eliminates product obsolescence and can reduce costs because developers can quickly and easily upgrade their designs to target the latest low-cost device. A general misconception that we see in the automotive design community is that FPGAs are too expensive for production. Five years ago, one million system gates cost around \$45. Today, these same one million system gate devices sell for less than \$10, with the smaller 100 K system gate designs selling at less than \$3, allowing for massive integration of multiple components into a single device. It is now possible to take an FPGA into full production and achieve the system cost targets required by the automotive market.

FLEXIBILITY AND RELIABILITY

The programmable nature of PLDs offers yet another level of advantage--in-vehicle programming and reprogramming. Device in-vehicle programming enables algorithms and functions to be upgraded even after product deployment. Since current telematics and video image recognition systems are in the early stages of research and development, the ability to make in-field upgrades can be a crucial asset. As technology--such as image-processing algorithms--improves over time, hardware upgrades can be accomplished in a matter of minutes without having to redesign an ASSP or lay out a new board.

For example, in instrument cluster and center stack display designs, low-voltage differential signaling (LVDS) transceivers have given automotive designers low-noise, high-speed signaling interfaces needed to implement flat-panel display (FPD) applications. Recently, various display manufac-

turers have adopted reduced swing differential signaling (RSDS) interfaces. This new signaling technology comes with a number of benefits over LVDS, including lower dynamic power consumption, further reduction in radiated EMI, reduced bus widths, high noise rejection, and high throughput. Again, the dynamic nature of PLDs gives developers the choice advantage. PLDs support a multitude of I/O signaling standards, giving developers the option to incorporate newly adopted technologies, such as RSDS, into their design.

On the reliability side of automotive design, there are many elements to consider. While ISO-TS16949 certification is a given in the market, a designer needs to take a deeper look. Many companies are using third-party subcontractors for their production. Designers must ensure the supplier itself is certified. If not, the supplier does not have its design and operational flows certified to the industry standard. In automotive telematics applications, AEC-Q100 automotive IC stress test qualifications and production part approval process (PPAP) documentation is also mandatory.

Back on the technical side, using PLDs will also improve reliability. Although LVDS transmitter and receiver pairs are readily available on the market, employing PLDs allow developers to integrate the transceiver onto a single device. PLDs not only offer various integrated signaling capabilities, but they also integrate source and termination resistors. By eliminating the multitude of discrete components, the designer achieves a reduced component count, resulting in a simplified PCB and a far more reliable signaling structure. The end result is a more cost-effec-

tive and reliable system.

In addition, they also provide the capability to contain an entire system on a single programmable device, including the processor. By placing an entire design on a single chip, designers can reduce the number of components on the board and their relative connections resulting in a scalable, portable and reliable system. Color temperature, for instance, is one of the many image enhancement issues facing vehicular display developers. Different regions around the world require different color temperature preferences. By using a PLD to create a scalable solution for color temperature adjustment, this solution can be leveraged across multiple geographies to support multiple display types, with minor adjustments toward the geographically preferred color temperature setting. Platform scalability and design reliability stay intact, while taking advantage of cost savings.

Most PLDs have built-in clock conditioning for duty cycle correction and clock managers that allow clock manipulations. The clock managers are placed on

internal dedicated, low-skew lines enabling precise, global clock signals. Such clocks enable a complete solution for high-speed clock designs, such as those needed for image processing. De-skewed internal and external clocks eliminate clock distribution delays and provide high-resolution phase shifting. These clocks also have flexible frequency synthesis, generating clock frequencies equal to a fractional or integer multiple of the input clock frequency. Dependable clock management systems are useful for timing and control circuits to meet growing display requirements.

Image scaling needs can also be addressed by PLDs. Take real-time image resizing, for example. The line buffers and coefficient banks can be implemented using block RAMs. Everything else including the vertical and horizontal multipliers, the adder tree, the sequencers and control, can be implemented using basic logic structures within the PLD. There is also no intermediate buffering necessary between the vertical and horizontal multipliers, therefore, there is no frame latency.

CONVENTIONAL DSP VS. FPGA

Many automotive telematics applications today demand high-performance video and image processing. PLDs have a number of features that make them ideal for handling applications, such as navigation systems and rear-seat entertainment/video. For example, distributed RAM on FPGAs is used to store DSP coefficients and finite impulse response (FIR) filters offering high memory bandwidth. Dual-port block RAMs are available for optimized data buffering and storage and for applications like fast Fourier transforms (FFTs). PLDs can also perform billions of multiply-accumulates (MACs) per second; using MACs built by embedded multipliers and accumulators. The large number of multipliers found in PLDs can also be used to create parallel multiplier arrays that support complex high-performance DSP tasks, where conventional DSPs are limited to serial processing (Figure 1). Embedded SRL16s, made up of registers and look-up tables (LUTs), enables highly efficient implementations of multichannel data paths. They can also dramatically increase FPGA compute density by enabling the construction of efficient time-division multiplexed (TDM) hardware structures.

Simply by using PLDs, developers can leverage its flexible architecture and take advantage of distributed DSP resources, such as LUTs, registers, multipliers and memory. With distributed DSP resources throughout the device, segmented routing, and component usage, FPGAs allow algorithms to be optimally implemented in the device. For example, designers can size an array to suit the exact calculation requirements, ideal for performing calculations on images. Calcula-

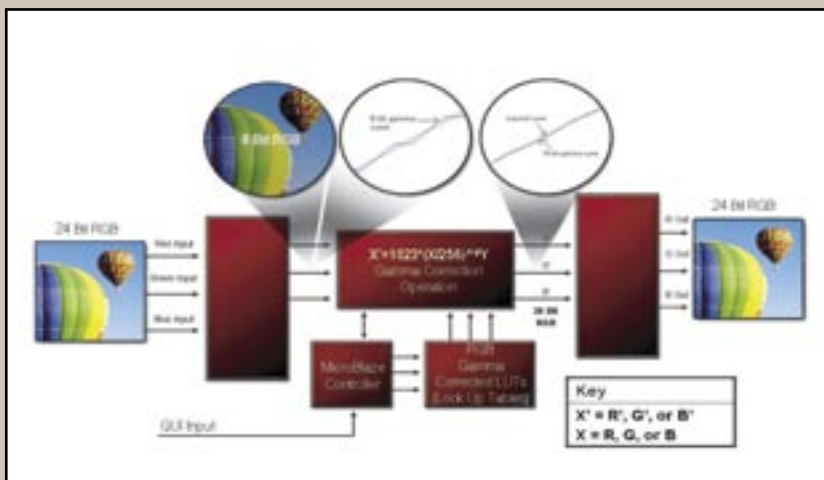


Figure 2. Gamma correction using FPGA.

tions can be performed on clusters of pixels, such as discrete cosine transform (DCT) blocks concurrently with other blocks in the picture, instead of having to scan the entire picture sequentially. And because processing can now be done in real time, less memory is needed for buffering pixel values when using PLDs.

Although conventional programmable DSPs can address a wide range of applications, it has its limitations. For instance, conventional DSPs are bound by their architecture, with fixed data widths and limited MAC units, therefore, limiting its data throughput by using serial processing. This forces the system to operate at high clock frequencies to increase data throughput creating yet another set of challenges. At the same time, it takes multiple DSPs to meet bandwidth needs, creating power and board space issues. By using a PLD, a designer can implement the custom solutions required to address higher-performance, high-quality, real-time display challenges. PLDs, with their flexible architecture and DSP resources, support serial and parallel processing. By opting to use parallel processing, a system has the potential to maximize its data throughput with a single clock cycle. Again, the designer can size the array to suit specific processing needs.

The issues that are usually addressed through custom, discrete ASICs, ASSPs, or graphic processors find their resolution in PLDs. For example, an image enhancement application for DSP is found in gamma correction needs for high-resolution LCD monitors. Gamma correction controls the overall brightness of an image. It can also affect the hue of a particular color representation, affecting the ratios

of red to green to blue. All graphic sources assume that display devices have a non-linear luminance input to output function, called the gamma function where $V_{out} = V_{in}$, where typically ranges from 2.2 to 2.8. If this discrepancy is not corrected the output display will have a pale appearance with little color saturation. In PLDs, gamma correction in the RGB space is generally done by dynamically updating an LUT to display the proper response at the output. When comparing an eight-bit vs. a 10-bit LUT approximation, it is apparent that the 10-bit resolution is a better approximation to the ideal gamma curve.

The formula for this approxima-

tion using 10-bit LUTs is:

$X' = 1023 * (X/256)^{(1/\gamma)}$ where $X' = R', G',$ or B' the 10-bit corrected output

$X = R, G,$ or B the 8-bit uncorrected input for gamma

Note: If the calculations yield a fractional result, normal rounding rules apply.

The gamma corrected 30-bit R'G'B' output needs to be fed through an image dithering engine that will find the closest color approximation for a 24-bit RGB output to the display device. There are several algorithms for dithering pictures. By using PLDs, developers can compare several algorithms quickly to determine which will meet their application

requirements. Dithering algorithms can also be quickly and easily changed by simply reconfiguring the PLD with the algorithm change in the source code.

The color temperature corrector is the feedback device that will dynamically change the input RGB values depending on the color response of the output. The values of the RGB outputs are compared to the color temperature of the blackbody

diagram of one such system. Running from DVD or HDD through the ATAPI interface, or pulling data from a USB interface, the MicroBlaze 32-bit soft microprocessor is used to run the rear-seat entertainment application code, DTCP AKE, disc navigation and file systems. Beyond that, it can also run the MOST DVD FBlock decode and net services. Data run through the ATAPI interface is connected to IP

to the TFT displays. If there is a secondary audio/video auxiliary input, two video processing and display controller IP blocks can be integrated into the FPGA to run multiple TFT displays, and each I/O port has the capacity to drive a separate standard.

The automotive industry is facing one of the most exciting and challenging times in its history. New modules are being implemented that include new and rapidly changing protocols, including some coming from the quickly evolving consumer market. Tougher schedule constraints are making it more difficult to maintain the high quality and reliability requirements of the automotive industry. Flexible and platform scalable integration at the system level is becoming a necessity in order to hit the low OEM module cost targets.

Today's PLDs have become a viable alternative to fixed logic devices. PLD suppliers are demonstrating their commitment to serving the automotive market by offering temperature-tolerant packaging of -40 °C to +125 °C, and striving to meet the stringent requirements of the automotive industry including ISO TS16949 certification, AEC-Q100 qualification flow and the PPAP. This allows automotive engineers to meet their challenging design goals with complete confidence in component quality and performance, while providing the ability to respond quickly to constantly changing automotive and multimedia standards and protocols. ■

ABOUT THE AUTHORS

Kevin Tanaka is the automotive marketing and product planning manager at Xilinx Inc. Cindy Kao is vertical marketing engineer at Xilinx.

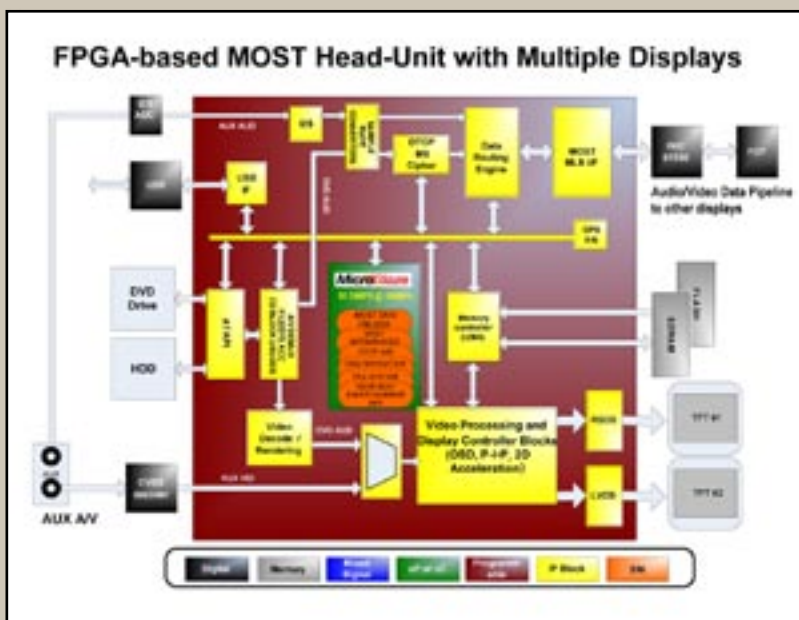


Figure 3. FPGA-based MOST head unit with multiple displays.

radiator to determine the ideal color temperature outputs on a dynamic basis. This can all be done in a single PLD as shown in Figure 2.

DRIVING MULTIPLE DISPLAYS

An example of the ability of an FPGA to drive multiple displays is especially evident in the mobile multimedia and rear-seat entertainment market. This is an area where all tier one engineers are looking for a flexible solution to handle multiple display standards and generate a system quickly. Figure 3 is a block

blocks implemented to do audio and video demux, along with CD block decode and can be routed to either the video decode and rendering block, run directly to a sample rate converter, through the DTCP cipher and out the MOST MLB interface, or dropped onto the OPB arbiter. On the video decode and rendering side, data is streamed into the video processing and display controller IP block where OSD, P-I-P and 2-D acceleration processes are completed before being sent out either RSDS or LVDS I/O interface directly

