



Solving Today's Interface Challenges With Ultra-Low-Density FPGA Bridging Solutions

A Lattice Semiconductor White Paper
August 2013

Lattice Semiconductor
5555 Northeast Moore Ct.
Hillsboro, Oregon 97124 USA
Telephone: (503) 268-8000
www.latticesemi.com

Designing Optimized Interface Bridging Solutions

Today's embedded system designers face an unprecedented challenge from an I/O perspective. As system complexity rises, they are increasingly asked to address a multitude of potential I/O options. These options can range from interfacing one industry bus to another, to connecting new and higher performance sensors with mature application processors. Moreover, this problem is pervasive across all markets from high volume consumer applications to the latest industrial, scientific and medical systems.

To be sure, few system architectures can deliver the tremendous flexibility needed to scale with today's rapidly changing I/O requirements, particularly where low power and low cost are concerned. To address this need, designers are implementing a wide variety of interface bridging solutions that allow them to transfer data across protocols and, in the process, expand system functionality. The challenge is determining how to most efficiently implement these new bridging solutions without violating system power, footprint and cost requirements. This paper looks at potential solutions and examines how designers can tackle the interface challenge by implementing highly optimized bridging solutions in ultra-low density (ULD), low power field programmable gate arrays (FPGAs) that combine the flexibility of a programmable platform with high performance at low power.

Wide Application Range

Interface bridging applications vary widely depending upon performance requirements. At the lower end of the performance spectrum designs often feature a variety of interfaces running under 50 MHz including UART, I2C, I2S, SPI, USB, Microwire, SoundWire, memory cards and PWM to name just a few. These bridging applications can take the form of converting the signal from a single I2C interface to many I2C interfaces to support multiple Small Form-factor Pluggable (SFP) transceivers in a communications subsystem, or converting from I2C to PWM for LED backlighting applications.

Of course, bridging requirements grow increasingly complex as clock rates rise. Designers working in markets from image processing and communications to state-of-the-art scientific, industrial and automotive fields, for example, need to find new ways to translate data from one protocol to another in a cost-effective manner. Typical applications might include the integration of new higher speed image sensors into legacy systems or the migration from one media standard to another. Ideal solutions are those with a proven track record of reliability and cost-effectiveness in highly competitive, high volume applications.

Design Tradeoffs

To address today's interface challenges designers need a highly flexible, quick turnaround solution capable of supporting higher data rates and specialized protocols without imposing long development cycles, increasing power consumption or driving up cost. Unfortunately most current design options fall short of those goals. For example, the available I/O's on most microcontrollers (MCUs) are speed limited. However, many of the higher speed buses mentioned above use differential signaling and a serial interface which, running as fast as 500 Mbps, far exceeds the capabilities of popular MCUs.

Moreover these high-speed buses typically employ highly unique commands and sync codes that most MCUs cannot support.

While many commonly-used MCUs will support lower-speed interfaces such as I2C, SPI and PWM, an MCU-based bridge would require a substantial investment in code development. Furthermore, as an interrupt-based system, the typical MCU operates at a distinct disadvantage. When an interrupt occurs it enters a queue that, depending on what other tasks the processor is executing, runs the risk of a latency that could undermine data accuracy. Any MCU operating as a bridge cannot allow an interrupt to stop the bridge function. Accordingly, the only option is to dedicate the entire MCU to perform the bridging task which, in most cases, will significantly increase system footprint, cost and power dissipation.

Dedicated ASICs offer a better match for bridging applications from a performance standpoint. But their long development cycles and high NREs price them out of the bridging market. ASSPs deliver a more attractive cost structure, but few ASSP solutions exist because manufacturers cannot justify their investment for non-standard solutions. Ultimately the relative inflexibility of any fixed-function solution renders most ASSPs too narrowly-focused for many bridging applications.

Ultra Low Density FPGAs

FPGAs have traditionally been viewed as too large, expensive and power-hungry for these types of bridging applications. However, a new class of ultra-low density (ULD) devices developed by Lattice Semiconductor is rapidly altering perceptions and opening new opportunities for building optimized bridging solutions. Leveraging recent technological advances aimed at reducing power consumption, footprint and cost, ULD FPGAs offer designers a highly unique set of advantages that traditional bridging solutions cannot match.

First and foremost is design flexibility. ULD FPGAs from Lattice Semiconductor, for example, offer a wide range of I/O counts from 21 to over 250. Since each I/O on the FPGA is programmable, designers no longer have to juggle interface requirements to meet the limitations of their applications processor or ASSP. With ULDs they can simply reprogram their design to support new interface bridging needs by implementing any form of interface on the FPGA's input and output. The high degree of flexibility allows designers to quickly and easily swap out interfaces as a design evolves or interface standards change without altering code in the applications processor. As designers migrate to non-standardized interfaces that require highly specialized clock rates and protocols, the ability to reprogram an interface and reallocate logic resources to meet changing needs offers embedded system designers an unmatched advantage.

At the same time, ULD FPGA approach to bridging can be implemented without paying a penalty in power consumption. Over the past two decades typical FPGA static power has dropped dramatically from 0.5W to under 50 μ W today. Accordingly, an FPGA-based bridge operates at an order of magnitude less power than a typical processor. Yet the power savings extend well beyond that. By implementing a bridge in real-time logic on an FPGA instead of an interrupt-driven processor, designers no longer need to keep the power-hungry apps processor operational at all times. The FPGA can respond instantly to

incoming data and effectively relieve the applications processor from time-consuming bridge management and control operations.


Finally, ULD FPGAs offer a reliability track record proven in the most competitive consumer applications. Widely employed in leading handset designs, they have demonstrated the uncompromising level of quality and the cost-effectiveness only high volume consumer markets demand.

Application Examples

The following design examples illustrate how development teams can use the inherent advantages of ULD FPGAs to meet their bridging needs.

IR to I2C

The first example is an IR to I2C bridge depicted in the block diagram below. This bridge uses a Lattice Semiconductor iCE40LP1K FPGA in a compact 36-ball wafer-level Chip Scale Package (WLCSP) to interface a Vishay TSOP34338 or TSOP98200 IR receiver module to the Qualcomm APQ8060-based development system. The APQ8060 does not have interface pins for an IR interface, thus the iCE40LP1K is ideal to perform the bus conversion from IR to I2C. Typical static current on the iCE40LP1K is only 40 μ Amps. There exists enough resources in this bridge design for additional pins to drive two IR transmitters if needed. Different configurations of the design will support different remote controls or IR receivers.




IR-to-I2C Bridge using an iCE40 FPGA from Lattice Semiconductor

Dual Image Sensor Bridge

A growing number of applications such as 3D stereoscopic video, 3D analytics for security or surveillance, and black box car driver recorders require multiple image sensors. While many of today's commonly-used image signal processors (ISPs) feature only a single image sensor port, they offer the performance capability required to process the data from two sensors if it can be correctly integrated. To expand the ISPs limited port configuration, designers use a dual sensor interface bridge.

The design illustrated below depicts a typical dual image sensor bridge for a black box car driver recorder. These systems typically use two cameras, one pointed out the front windshield and a second pointed at the driver. The bridge is implemented in a Lattice Semiconductor Mach X02-1200 ULD FPGA. It synchronizes and manages the two image sensors and outputs the data in a format the ISP can accept. The X02-1200 FPGA features up to 1200 LUTs and 108 I/O. It also adds hardened blocks for I2C, SPI and a timer/counter. The dual sensor bridge can support two image sensors up to 720p resolution each. The output format for the ISP is either a top/bottom or left/right configuration.




Dual image sensor bridge using a MachX02 or iCE40 FPGA

GMII-to-RGMII Bridge

Another common bridging function is keeping up with evolving standards. The Gigabit Media independent interface (GMII) is a widely-used Ethernet interface standard. However, newer designs are migrating to the Reduced Gigabit Media independent interface (RGMII) which was developed to reduce cost by driving down the number of pins from 22 to 12.

The design below illustrates the implementation of a bi-directional GMII to RGMII bridge using an FPGA from Lattice Semiconductor. The bridge operates at > 125 MHz and uses HSTL I/Os without any additional buffers. The pins are divided into a GMII and a RGMII side. Data is transferred at double data

rate which allows transfer on both the rising and falling clock edges. The FPGA features I/O shift registers (IOSR) for each group of programmable I/Os PIOs that are configured to work together and transfer data on both clock edges.




GMII-to-RGMII bridge

MIPI CSI-2-to-CMOS Parallel Sensor Bridge

As user demand for integrated cameras in smart phones has soared, the MIPI Camera Serial Interface-2 (CSI-2) is becoming an increasingly popular solution for imaging applications. The specification defines a high speed, low power interface between a peripheral camera and a host processor. Given their high performance and low cost, CSI-2-compliant image sensors are now widely used in embedded applications. Most ISPs in embedded use today, however, feature a traditional CMOS parallel interface that requires a large number of interconnects and consumes a relatively large amount of power. As such, these ISPs cannot connect directly to an image sensor that features a CSI-2 interface.

The design below depicts a MIPI CSI-2- to-parallel bus image sensor bridge. It supports CSI-2 high speed differential signaling from one to four data lanes running at up to 750 Mbps. Implemented in a Lattice Semiconductor MachXO2 FPGA, the bridge supports 200mV common mode voltage in High Speed mode via the FPGA’s differential input pads. It also provides parallel sensor output of 6 to 24 bits RAW, RGB and YUV. While the bridge is typically synthesized in a single CSI-2 format to minimize FPGA size, it can support multiple CSI-2 formats for “on-the-fly” switching by adding multiple instantiations of the MIPI CSI-2 Serial2Parallel NGO in each desired format.



MIPI CSI-2- to- Parallel bus image sensor bridge using MachXO2 FPGA.

Image Sensor Extender

In a number of applications, such as pipe inspection, medical probes and some aftermarket automotive cameras, an image sensor must be located far from the ISP. To address this requirement, Lattice Semiconductor developed a sensor extender that leverages two MachXO2 FPGAs, one positioned near the sensor to serialize the signal and a second near the ISP to de-serialize the signal, with a single 5e/6 cable in between.

The image sensor can operate at 720p60 or 1080p30 at up to 8 meters distance with well-matched pairs. The interface to the ISP after de-serializing can be parallel or serial. Parallel interfaces can be configured for a variety of LVCMOS levels including 1.8V, 2.5V or 3.3V. The extender can be used with two or more sensors. It also supports interfaces to CSI-2, HiSPi, parallel and sub-LVDS image sensors.




Image sensor extender using two Lattice XO2 FPGAs and a Cat 5e/6 cable.

Conclusion

As long as designers continue to add new functionality to their systems, I/O capability will continue to grow. Bridging solutions provide a highly practical way to match rapidly expanding I/O capability with lagging system architectures. As a highly flexible development platform available with a diverse range of I/O and throughput capability, ULD FPGAs offer designers a unique opportunity to build optimized, power efficient bridging solutions for their next generation systems.