
Power-Optimized Solutions for Telecom Applications

Introduction

Some telecommunications (telecom) carriers now target a 20 percent power reduction per year on the equipment they deploy. This is a response to increasing power consumption in central offices, which is due to higher bandwidth capacities, increasing linecard port densities, more intelligent processing requirements, and more complex chip implementations. Higher power consumption equals higher operating costs for telecom providers, and leads to more complex engineering challenges for equipment suppliers dealing with thermal management.

Semiconductors are a significant part of this power problem. The answer is to re-think chip implementation and delivery in order to meet the aggressive power-reduction goals set by carriers. The key enablers are the latest generation of FPGAs. Where FPGAs historically were associated with high power consumption, they are now best suited to solve the power consumption problem. By adopting the latest semiconductor manufacturing processes based on 40-nm geometries along with innovative approaches to optimize these complex devices, designers can integrate more functionality on a single device. This allows for a reduction in overall power consumption as well as the ability to reduce the power per comparable function in future process nodes.

One company taking advantage of power reduction is TPACK. TPACK provides Carrier-Ethernet chip solutions based on Altera® Stratix® IV FPGAs to the majority of the world's largest telecom system vendors. The combination of Altera's high performance and power optimizations with TPACK's expertise in providing complex and highly integrated devices results in a compelling proposition to system vendors. Not only do they continue to meet demands for higher bandwidth capacity and more intelligent processing, but they do so using chip solutions that provide lower power consumption.

In addition, the chip solutions provided by TPACK can be ported to the latest FPGAs, reducing power even further. The net result is system implementations that offer lower power consumption now, and will continue to meet demands for power consumption reduction in the coming years.

The Power Issue Is Getting Hotter

Power consumption has always been a major consideration when building systems and chips, so why the increased focus now? The issue is not just an issue of global warming or adopting the latest trend. Carrier operating costs have grown as more sophisticated cooling systems are needed and more electricity is consumed, thereby leading to higher energy costs. This is only compounded by the fact that energy costs have risen considerably of late.

According to EnergyChoices, a consumer service tracking energy prices in the UK, energy costs in 2008 rose by 66 percent for electricity and 60 percent for gas. This is not a recent phenomenon either. Data from the Energy Information Administration in the U.S. states that the average retail price of electricity increased 27 percent from 2000 to 2006. It is therefore likely that energy prices will continue to grow.

These statistics have prompted carriers to make reducing energy consumption a critical focus area. For example, in June of this year, Verizon became the first carrier to set energy-efficiency goals for network, data center, and customer equipment, aiming for 20 percent more efficiency by 2009. Verizon, BT, and AT&T have defined strict requirements for power efficiency of various types of equipment. Other carriers are following suit by including energy efficiency as part of their requirements in the purchase process.

The Power Challenge

Meeting these strict power efficiency requirements is a major challenge for system developers. Reducing power consumption by 20 percent is not trivial as power consumption has increased over the last few years.

System Perspective

One major source of system power consumption is the semiconductor chips used to process traffic and the associated memories. For Internet protocol (IP) routing, which requires large table lookups and large memories, four to eight ternary content-addressable memories (TCAMs) are required. Because TCAMs typically consume 15W, a total power consumption of 120W is possible.

System developers address this issue by using more dedicated processors and performing more operations at layer 2 and 2.5, using Ethernet and MPLS to avoid TCAMs. DRAM and SRAM memories can be used instead, typically consuming only 2W each. The processors that take advantage of low-power memories then become the major source of power consumption, based on the complexity and bandwidth of the chips involved.

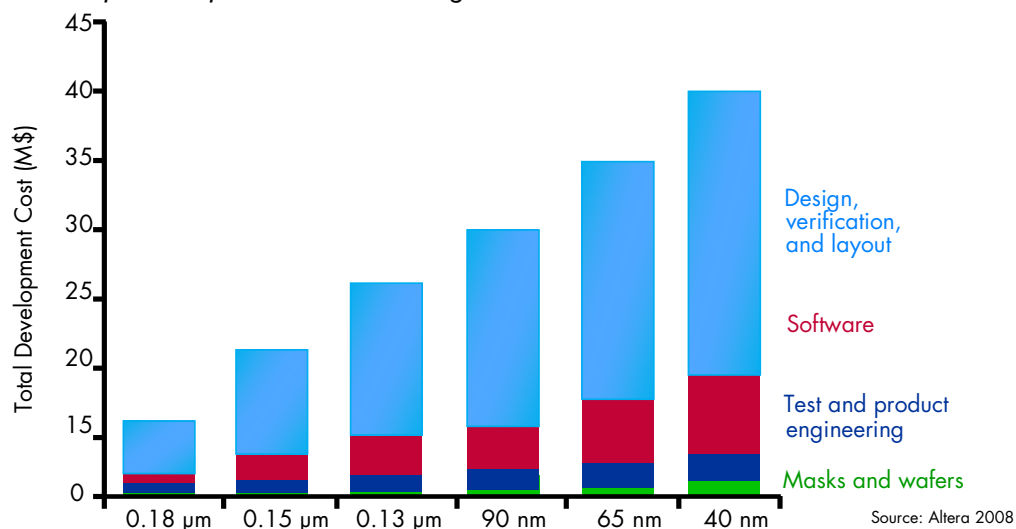
For example, take Broadcom's range of control processors used in routers. For the BCM1122 single-core processor with one 10/100/1000 Ethernet interface and one 10/100 Ethernet interface, the power dissipation is 4W. However, this doubles as the number of cores doubles in support of extra interfaces. Therefore, the dual-core BCM1250 consumes 8 to 10W and the quad-core BCM1455, with three Gigabit Ethernet interfaces, consumes 19W.

Network processing units (NPU) for packet processing also consume a great deal of power as more complexity is added. The EZChip 10-Gbps NP-1c consumes 15W, the 20-Gbps NP-2 consumes 17W, and the 30-Gbps NP-3 consumes 20W. If each extra 10 Gbps of switching capacity adds 2 to 3W of power consumption, even taking into account moving to lower geometries, what does this mean for fast-approaching 100-Gbps processing requirements?

Semiconductor Economics

Having established that the complexity and capacity of chip solutions grow as they progress to 100-Gbps systems, the solution for addressing power consumption is to pursue the latest advances in semiconductor manufacturing processes aggressively. This enables higher density, higher bandwidth chips, and systems with relatively lower power consumption. However, the cost of developing chips at lower geometries increases exponentially with each step down in process geometry (Figure 1). This makes a challenging business case for specialised chips addressing high-end telecom systems.

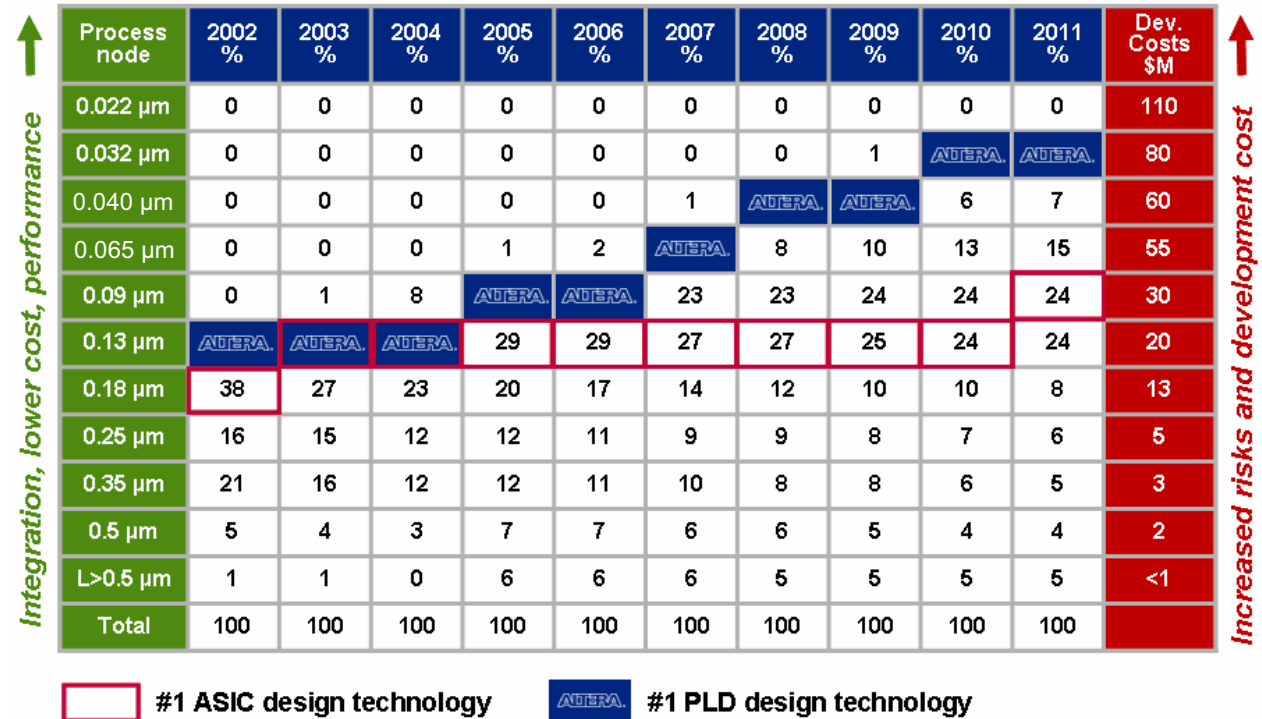
Figure 1. Cost of Chip Development for Decreasing Process Nodes



The number of systems and chips required in high-end telecom systems are measured in the thousands, not the millions normally associated with semiconductor manufacturing. The many competing protocols and approaches fragment this market even further, making it difficult to predict revenue streams for new chip developments with any degree of confidence or accuracy. This unpredictability has already discouraged many ASSP chip suppliers from migrating from 130-nm to 90-nm processes with very few making the move to 65-nm processes and beyond.

In contrast, many early adopters of the latest manufacturing processes are FPGAs (Figure 2). FPGAs serve a broad base of applications in many industry verticals, including communications, military, automotive, and consumer products. This provides a volume base that justifies investment in the latest semiconductor manufacturing processes.

Figure 2. ASIC Design Starts by Technology



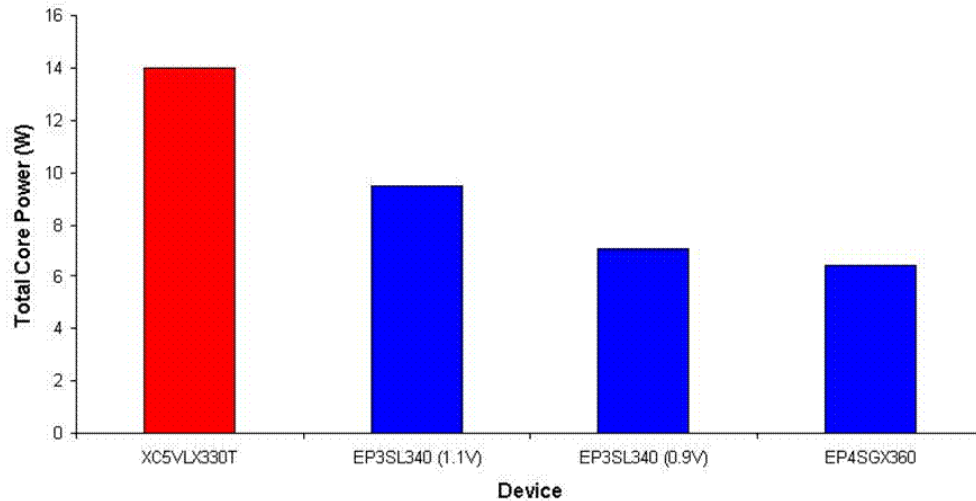
Source: Altera & Gartner Nov 07

Silicon Perspective

As semiconductor process geometries shrink, there are variables, such as reduced capacitance and voltage, that lead to natural decreases in dynamic power consumption. For example, compared to a 65-nm process node with the same architecture, the lower voltage (0.9V) of a 40-nm process results in a 33 percent reduction of dynamic power, and the 30 percent capacitance reduction further reduces dynamic power.

However, the reduction in dynamic power may be offset by significant increases in power due to other factors. For example, static power may increase dramatically on smaller geometries due to various sources of leakage current. In addition, dynamic power may increase due to the greater density and clock frequencies of the smaller geometries. So, building silicon at 40 nm and beyond without aggressively designing for power optimizations results in solutions that move the power curve in the wrong direction. Altera proactively built the 40-nm Stratix IV family with power optimization as a key objective, thereby reducing power by several watts compared to competing FPGA technologies (Figure 3).

Figure 3. Reducing Power Consumption by 50%




Addressing the Power Challenge

Tackling power consumption in leading-edge silicon technologies involves a variety of approaches, including process, architectural, and design optimizations. This section describes process optimizations, architectural optimizations, and design optimizations for developing power-optimized solutions. A key factor in delivering these power optimizations, called Programmable Power Technology, allows individual logic array blocks (LABs), memory, and digital signal processing (DSP) blocks to selectively turn on power savings based on specific design requirements.

Process Optimizations

Various techniques optimize Altera's 40-nm FPGAs for power, each with a different balance of benefits and drawbacks:

- Multiple-gate oxide thicknesses (triple oxide)
 - Trades static power for speed per transistor
- Multiple-threshold voltages
 - Trades static power for speed per transistor
- Low-k inter-metal dielectric
 - Reduces dynamic power, increases performance
- Super strained silicon
 - Increases electron and hole mobility by 30 percent
 - Balances between power and performance
- Copper interconnect
 - Increases performance, reduces IR drop

 For more information on these process optimizations, refer to Altera's white paper, *40-nm FPGA Power Management and Advantages*.

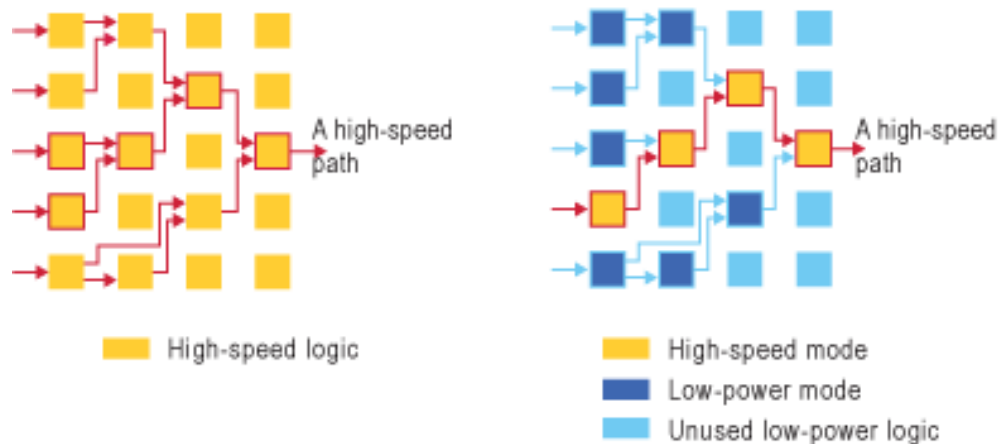
Architectural Optimizations

Telecom line cards often make scheduling decisions based on incoming packet traffic. This requires high-performance external memories to buffer the packets while the scheduling decisions are made. Stratix IV FPGAs offer dynamic on-chip termination (OCT) to reduce linecard power consumption. Dynamic OCT disables parallel termination on write operations, saving static power when writing the packets into memory.

Design Optimizations

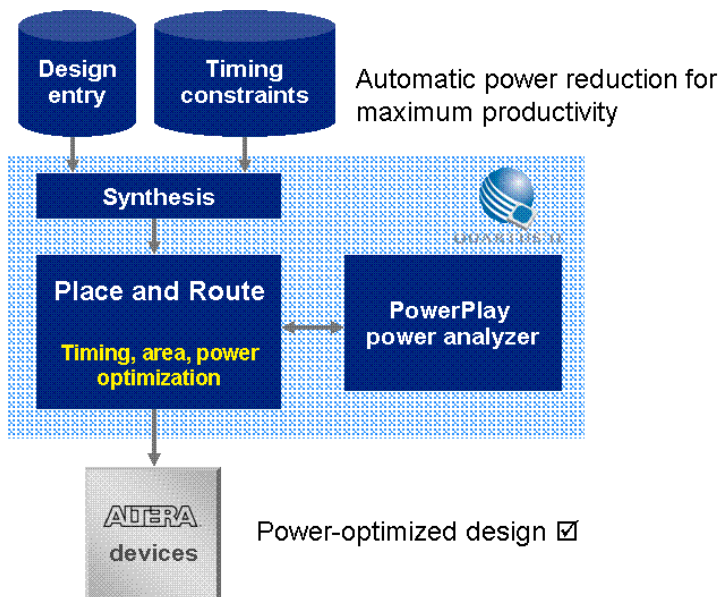
Programmable Power Technology enables every programmable LABs, DSP block, and memory block to deliver high speed or low power, depending on the design requirements. FPGAs not optimized for power contain blocks designed to run at only one speed—the highest possible speed—to support timing critical paths (as depicted by yellow blocks in Figure 4). Using Altera’s Programmable Power Technology, all LABs in the array, except those designated as timing critical, are set to low-power mode (as depicted by blue blocks in Figure 4). With only the timing-critical blocks set to high-speed mode, power dissipation is reduced substantially.

Figure 4. Standard FPGA Fabrics (left) Compared to Stratix III FPGA Fabric With Programmable Power Technology (right)



Another key area where Altera provides innovation is in the ability of Quartus® II development software’s synthesis and place-and-route engine to be power aware. This power reduction method is transparent to designers and enabled through simple compilation settings. The design engineer simply sets the timing constraints as part of the design entry process and synthesizes the design to meet performance. As shown in Figure 5, Altera and third-party tools automatically select the required performance for each piece of logic as well as minimizing power through power-aware placement, routing, and clocking.

Figure 5. Quartus II Software’s Synthesis and Place-and-Route Engine



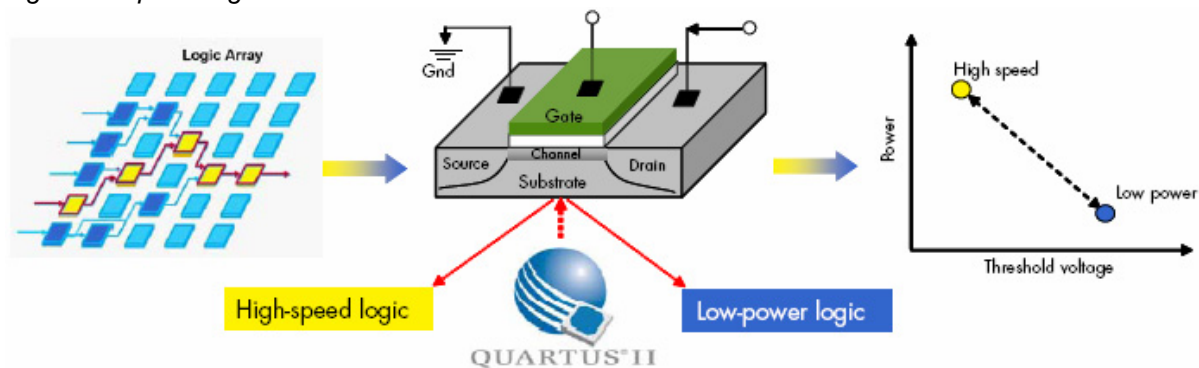
The resulting design meets the designer's minimum power requirements. The designer has the option to select low-effort or high-effort optimization. Selecting high effort offers the greatest power savings at the expense of longer compilation times. Results vary based on design and effort level selected. The goal of this feature is to reduce power without designer intervention while having minimal impact on design performance.

Power Optimization With Quartus II Development Software

Power optimization with Altera's Quartus II software occurs in three stages. Quartus II software first performs a "power-aware" synthesis. Power aware means that the software can minimize the number of RAM blocks accessed at each clock cycle, or re-arrange the design to eliminate high-toggling (or glitch-prone) logic.

After the power-aware synthesis, Quartus II software performs a power-aware place and route by routing signals to minimize capacitance or by creating power-efficient DSP block configurations. The PowerPlay power optimization option in Quartus II software guides the fitter to use extra effort to optimize the design for power by taking advantage of the power-specific architecture features. For example, Figure 6 shows how Stratix IV core-logic tiles operate in either high-speed mode or low-power mode. Using timing constraints, Quartus II software ensures the critical paths in a design are optimized for performance, while non-critical paths are optimized for power instead. This has a positive impact on the core static-power consumption, typically reducing it by up to 37 percent.

Figure 6. Optimizing a Critical Path



The assembler places unused areas of the FPGA in a low-power state, with additional power saved by minimizing toggling of unused logic. The PowerPlay power analyzer accurately estimates the design power using Altera's advanced modeling technology. The power analyzer generates a detailed power estimation analysis, which allows designers to manage and validate a power budget throughout the design cycle. It also includes a power optimization advisor that provides suggestions and information to minimize power consumption.

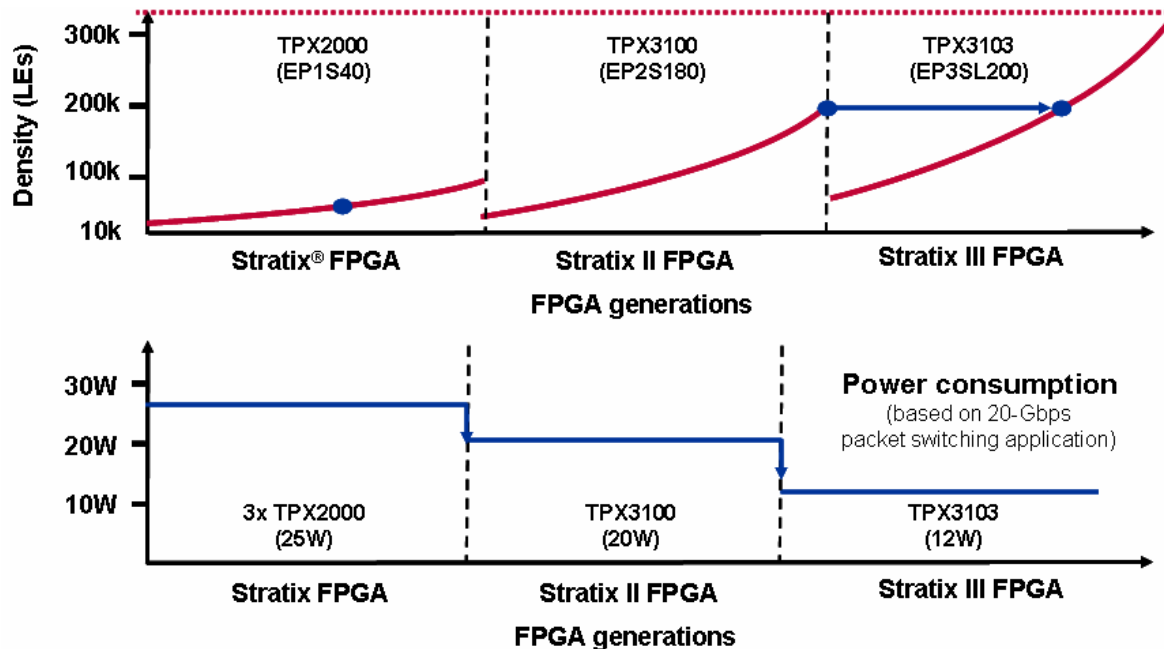
Leveraging FPGAs

By basing chip development on FPGAs rather than on foundry-based manufacturing, it is possible for an application-specific chip product to take advantage of the latest developments in chip manufacturing to provide power-efficient solutions. This approach, named SOFTSILICON by chip vendor TPACK, uses Stratix IV FPGAs to provide Carrier Ethernet packet-processing, traffic-management, and packet-mapping chip solutions to telecom system vendors.

 For more information on SOFTSILICON, see TPACK's white paper: [SOFTSILICON for Flexible Packet Transport](#).

TPACK uses the SOFTSILICON concept not only to enable the development of new chips with higher capacities, but also to reduce the power consumption of existing designs. As Figure 7 shows, TPACK's Carrier Packet Engines, which provide integrated Carrier Ethernet packet processing and traffic management, are based on various generations of Stratix series FPGAs manufactured using different geometries.

Figure 7. Reducing Power in Existing Chip Solutions



By moving from Stratix to Stratix II FPGAs, TPACK increased the capacity of the Carrier Packet Engine from 6 Gbps (TPX2000) to 20 Gbps (TPX3100), while decreasing the relative power-per-Gbps switching capacity. However, by porting this solution from Stratix II to Stratix III FPGAs (TPX3103), they reduced power consumption by as much as 40 percent.

The power-consumption advantages provided by the SOFTSILICON approach are two-fold:

- Provide more switching capacity at a relatively lower power per Gbps by adopting the latest FPGA platforms
- Reduce the power consumption of existing designs by porting to the latest FPGA platforms

These advantages help meet the stringent requirements for power reduction that carriers are demanding both now and in the future.

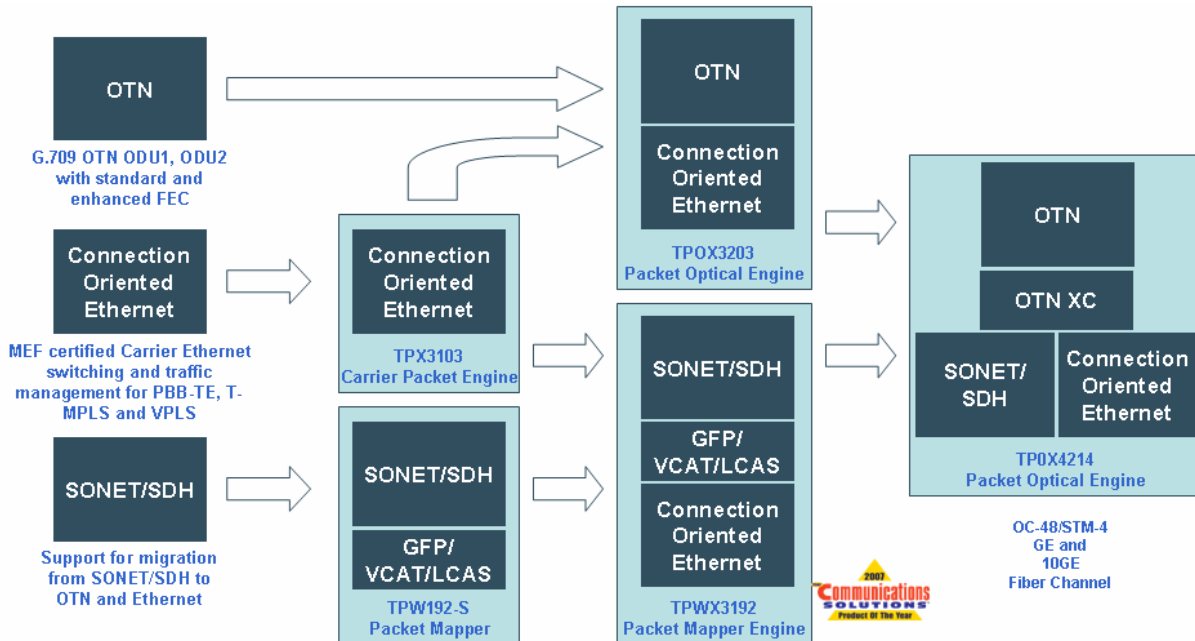
Leveraging SOFTSILICON

While SOFTSILICON takes advantage of the FPGA benefits described earlier, this is only part of the story. FPGAs provide an alternative chip-development approach to standard fabless models, which helps specialized chip vendors to keep up with system vendor and carrier demands. There is still a need for specialist chip suppliers such as TPACK with expertise in developing highly integrated, complex chip solutions that can deliver the advanced features required by the market.

FPGAs traditionally are used to implement relatively few functions (as an add-on to ASSPs) or as “glue-logic” to connect two incompatible ASSPs. It is only in the last few years that more complex chip solutions have been attempted on FPGAs, with TPACK one of the very first companies to accomplish this. With SOFTSILICON, TPACK provides a real alternative to ASSPs that often provides better performance, especially in relation to power consumption. In this respect, integration also plays a part. The ability to integrate more into a single chip ensures lower real-estate requirements and lower power consumption, as well as the flexibility to update the solution and fix errors quickly if needed. It is important to note that the disadvantage of ASSP integration is that the risk of errors increases with complexity, while with FPGAs, these errors can be corrected quickly, minimizing the apparent risks of integration.

Figure 8 shows the evolution of TPACK chip solutions in addressing the emerging Packet Optical Transport Network (P-OTN), which integrates connection-oriented Ethernet switching and traffic management with SONET/SDH and Optical Transport Network (OTN) technology. These solutions are based on Stratix III and Stratix IV FPGAs, which allow highly integrated devices. They leverage more than seven years of work on Carrier Ethernet switching, traffic management, and packet mapping, as well as TPACK’s proven experience in providing highly integrated solutions.

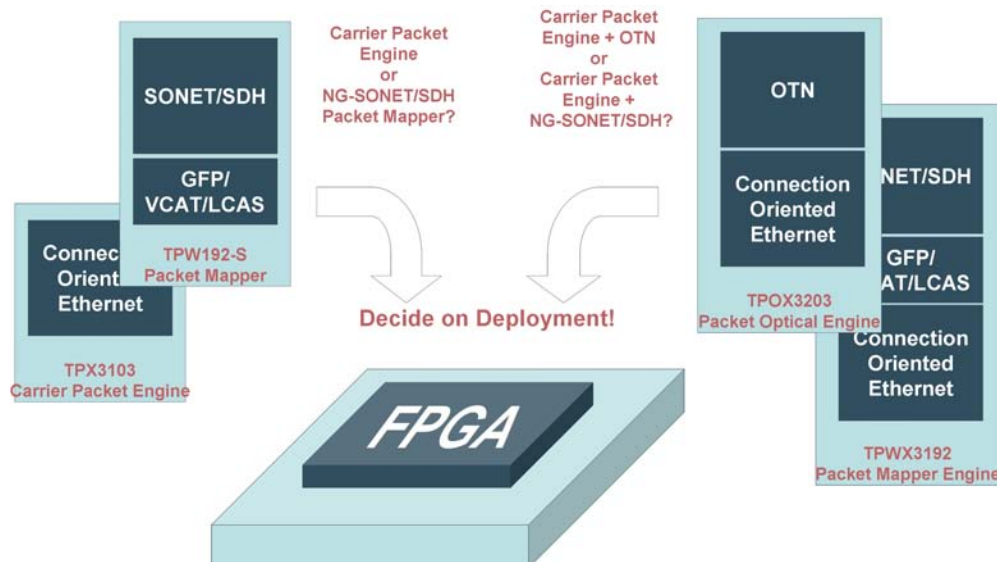
Figure 8. Integrated Packet Transport Solutions



Building Power-Efficient Systems Using Universal Linecards

From the discussion above, it is clear that there are power consumption advantages to using SOFTSILICON products in system design. However, SOFTSILICON provides even more power-saving potential if a universal linecard approach to system design is adopted. Figure 9 shows some of the possibilities that can be implemented.

Figure 9. One Chip Supporting Multiple Feature Sets



A universal linecard, or multi-function linecard, is based on SOFTSILICON and pluggable optics. This allows different interfaces and logic to be defined as needed, allowing a single linecard hardware design to be used for multiple applications. For example, in one application, the linecard supports Ethernet over NG-SONET/SDH packet mapping, while in another application it supports Carrier Ethernet switching and traffic management. All that is required is a linecard design prepared for the various options it is expected to support and a number of SOFTSILICON FPGA images to program the underlying FPGA platform during power-up.

Advantages of the Universal Linecard Approach

The universal linecard approach results in many system-development advantages, but from a power perspective, the most important advantage is the ability to reduce the number of defined functions on the card at deployment. In a typical chip development, all of the potential feature and interface options that system vendors and their carrier customers might require are designed into the chip. All possible options are included because if one is needed later and is missing, a costly re-spin is required to add it to the ASIC or ASSP. In practice, most of these features are not required by the carrier in any given deployment, but all must be present in the chip in order to ensure as broad an application as possible.

By adopting a universal linecard approach, only the features that are required for the particular deployment need to be loaded into the FPGA. This means fewer features, less logic, a smaller FPGA chip, and less power consumption. This approach requires managing a number of FPGA image variants, but has the advantage of allowing accommodation of carrier-specific requirements.

Conclusion

While the power-consumption reduction requirements of carriers seem daunting, meeting power optimization requirements is possible, and the solution is at hand. Altera's 40-nm Stratix IV FPGAs, TPACK's expertise with SOFTSILICON, and a universal linecard approach to system design provide all the ingredients needed to not only meet, but possibly exceed customer demands, while also enabling differentiation, responsiveness, and fast time to market with higher capacity packet-transport solutions.

Further Information

- EnergyChoices:
www.energychoices.co.uk
- Energy Information Administration:
www.eia.doe.gov
- *40-nm FPGA Power Management and Advantages:*
www.altera.com/literature/wp/wp-01059-stratix-iv-40nm-power-management.pdf
- *SOFTSILICON for Flexible Packet Transport:*
www.tpack.com/fileadmin/user_upload/Public_Attachment/SOFTSILICON_for_Flexible_Packet_Transport_w eb_v1.1.pdf

Acknowledgements

- Kevin Cackovic, Sr. Strategic Marketing Manager, Communications Business Unit, Altera Corporation
- Seyi Verma, Sr. High-End Technical Analysis Staff, High-End FPGA Products, Altera Corporation
- Daniel Joseph Barry, Director of Marketing, TPACK



101 Innovation Drive
San Jose, CA 95134
www.altera.com

Copyright © 2009 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.