

## Leveraging Cost-Optimized FPGAs to Deliver OTN Mapper Solutions

*In telecommunications transport infrastructure, Optical Transport Network (OTN) and Gigabit Ethernet (GbE) protocols are being combined to create Packet-Optical Transport Networks (P-OTNs). Telecom equipment manufacturers provide these new P-OTN systems and, naturally, are now looking to the component providers for silicon that supports these requirements. Products such as TPACK's TPO124, based on Altera's low-cost 40-nm Arria II GX FPGAs or HardCopy ASIC technology, allow innovative developers to offer flexible semiconductor products with lower cost and lower power than alternative solutions based on standard cell-based ASICs.*

### Introduction

Carriers around the globe are planning how to transition their existing telecommunications transport infrastructure from traditional SONET/SDH to a more cost-per-bit effective transport technology. The main drivers for this transition are the increased demand for bandwidth as well as the need to lower the cost per bit. Consumers of bandwidth expect higher and higher data rates on their fixed lines as well as our wireless connections. At the same time, however, they expect the price to stay the same or, better yet, decrease. This dilemma forces carriers to consider new and more cost-effective transport equipment than traditional SONET/SDH systems, especially for aggregation networks that aggregate data traffic from large amounts of users for onward transport in the core networks.

In some of these new transport systems, Optical Transport Network (OTN) replaces SONET/SDH as the transport protocol, and Gigabit Ethernet (GbE) replaces PDH E1/T1 (2 Mbps/1.5 Mbps) as the preferred access point into the aggregation network. As the majority of the data transported is packet based rather than the traditional circuit-switched telephony, these types of systems are often referred to as Packet-Optical Transport Networks (P-OTNs).

OTN is a standard defined by ITU-T, which also standardized SDH (based on SONET) twenty years ago. Currently, OTN is mainly used for point-to-point links of complete wavelengths on the fiber. However, OTN can provide connectivity at the sub-wavelength level via the Optical Data Unit (ODU) container. For example, a 10-Gbps optical wavelength (OTU2) may carry four ODU1 containers of 2.5 Gbps each. In April 2009, ITU-T approved amendment 3 to the G.709 standard (sec. 17.7.1), which defines a new ODU container, ODU0, that is half the size of ODU1 and can carry a GbE stream transparently.

Telecom equipment manufacturers such as Alcatel-Lucent, Huawei, Fujitsu, Tellabs, Nokia-Siemens-Networks, and Ericsson provide these new P-OTN systems and, naturally, are now looking to the component providers for silicon that supports the P-OTN requirements such as GbE mapping to ODU0- and ODU0/ODU1-level cross connects.

Infonetics forecasts that the worldwide P-OTN equipment market will double from around \$1.5 billion in 2008 to \$3 billion in 2011, despite the economic downturn. With a potentially high volume of P-OTN systems involved, the traditional approach to filling the OTN mapping solution need would be to develop a standard cell ASIC-based application-specific standard product (ASSP) for the telecom equipment manufacturers. However, more innovative implementation approaches can deliver the needed functionality, coupled with more optimized feature sets, in shorter time, and at a potentially lower cost. This white paper discusses one such implementation strategy, starting with an examination of the required feature set, then compares the implementation to a traditional standard cell ASIC-based ASSP approach.

### Requirements for OTN Mapping

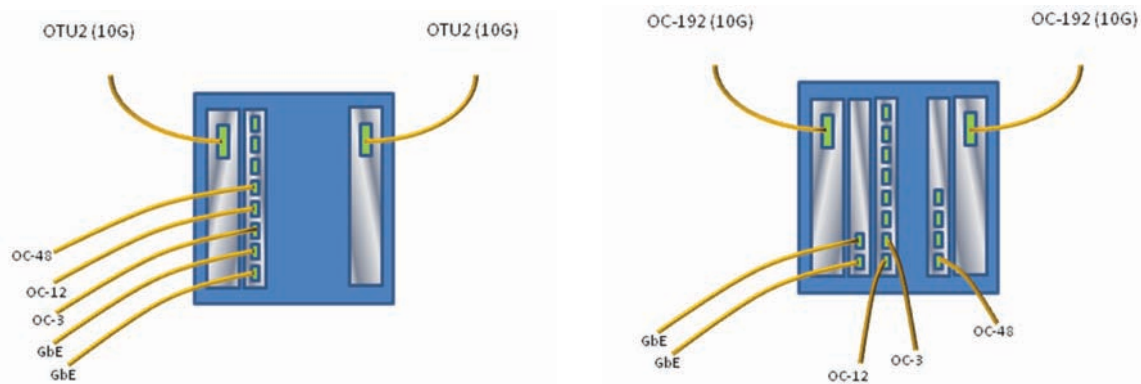
The key characteristics of the chip-level OTN mapping solutions are:

- Dynamically configured and reconfigured client ports to support any rate and any protocol
- Effective mapping of clients that are below 2.5 Gbps to OTN
- Cross-connection support at the ODU0/ODU1 level

Providing support for any rate and any protocol is one way to lower the capital expenses and operating expenses of P-OTN equipment. By doing so, any given hardware board can be configured to support any mix of protocols, for example OC-12, Fibre Channel, and GbE ports. With each network node in a carrier network supporting a different mix of protocols with a single piece of hardware, almost any access type needed by the carrier's customers at any given location is available. In addition, if the customer decides to upgrade from, for instance, OC-12 to GbE, the hardware just needs to be re-configured.

In contrast, SONET/SDH multi-service provisioning platforms (MSPPs) provide each type of client (for instance GbE and OC-12) with its own tributary card. Therefore, if a given site supports both GbE and OC-12, at least two tributary cards are required. [Figure 1](#) shows how a mix of OC-3, OC-12, OC-48, and GbE connections are supported by a traditional MSPP and new P-OTN equipment. This example illustrates that a P-OTN system requires much fewer tributary cards, thereby costing less to install and manage.

*Figure 1. Multi-Rate and Protocol Clients as Supported in P-OTN Equipment (left) and MSPPs (right)*



Until recently, the lowest OTN rate was 2.5 Gbps (ODU1). This rate does not match very well with the needs of typical OTN clients such as OC-3/OC-12 SONET (155 Mbps/622 Mbps) and GbE (1.25 Gbps). If any of these clients were mapped directly into an ODU1 container, between 50% and 90% of the bandwidth would be wasted. Therefore, it is desirable that a more effective sub-rate mapping be supported.

In addition, support for the newly standardized GbE to ODU0 mapping is crucial, because it enables the most important interface rate and protocol (GbE) to be transported in a fully transparent manner, with 100% bandwidth utilization inside OTN. In order to build cost-effective and operational aggregation networks, the system must be able to demultiplex to and cross-connect at lower rates. This is accomplished by building in ODU add-drop/cross-connect functionality in the P-OTN nodes.

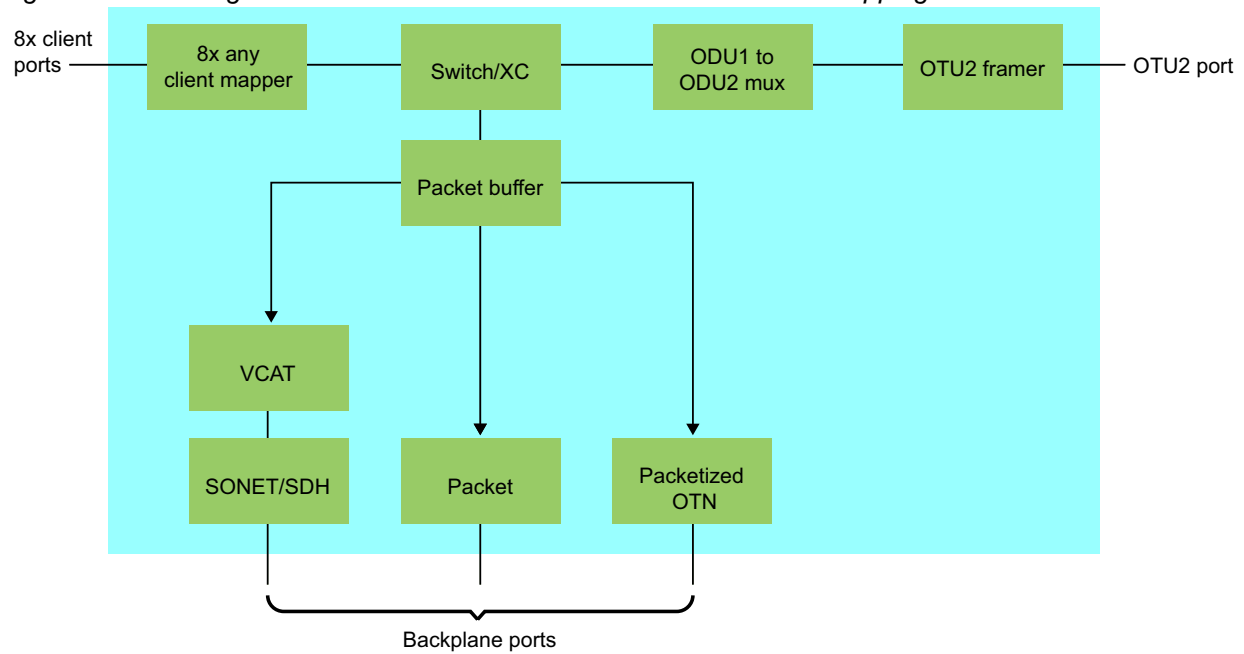
### Available Device Solutions for OTN Mapping

Until this year, the chip-level solutions for OTN mapping performed a straightforward mapping of typically 10-Gbps clients (such as OC-192 or 10 GbE) into an OTN transport container of the same speed, such as OTU2. Historically, framer chips with proprietary enhanced forward error correction (EFEC) algorithms and “home-grown” ASIC/FPGA solutions have dominated the market. However, none of these devices supports the sub-rate client ports or effectively maps the ODU0 described in the previous section. Therefore, new devices were introduced in 2009 to complement the simple 10-Gbps framer devices.

[Figure 2](#) shows a block diagram of one of the newer standard cell ASIC-based ASSPs for OTN mapping. It includes support for up to eight client ports, one OTU2 line port, and a range of backplane ports intended for connectivity towards a central packet or SONET/SDH switch. As with most ASSPs, the device in [Figure 2](#) is a superset of features intended to cover many different equipment architectures. This means that, in most applications, only a fraction of the device features are used, whereas the non-used parts are “dead” silicon that increase power consumption and cost. For instance, if the network equipment is a basic muxponder that multiplexes various lower rate clients into OTU2, none

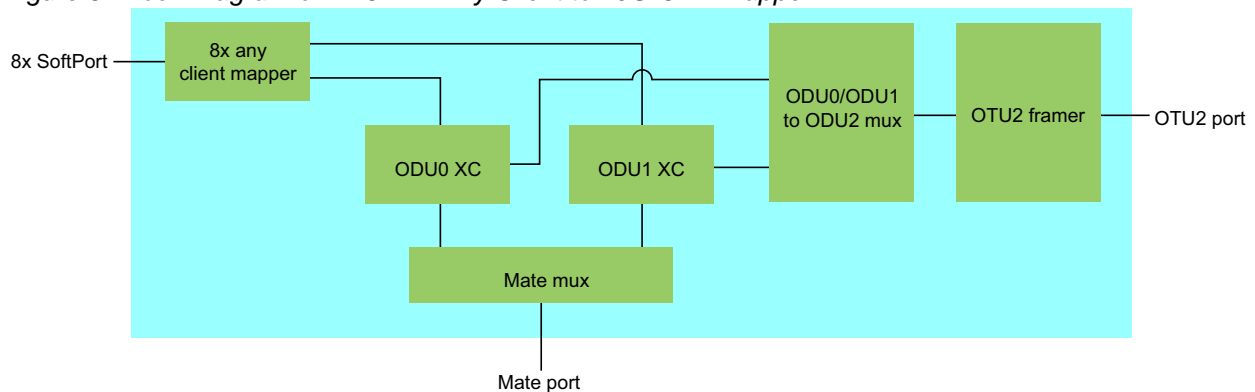
of the backplane ports are used. In addition, in larger OTN cross-connect equipment, a given chassis will have either a central packet switch or a central SONET/SDH switch. Many new P-OTN equipment products are built around a packet switch, which scales to higher capacities better than the older SONET/SDH switches. Therefore, it is unlikely that the large VCAT and SONET/SDH backplane port are used in most cases.

Figure 2. Block Diagram of Standard Cell ASIC-Based ASSP for OTN Mapping



Another recent OTN mapper device is the TPO124 from TPACK. Figure 3 shows a block diagram of the TPO124, which is a 2.5-10G OTN mapper. This block diagram illustrates support for effective sub-rate mapping, ODU add-drop/cross-connect functionality, and support for GbE mapping to ODU0. On the left are shown eight SERDES-based client ports, which can be dynamically configured to support any rate and protocol. The clients are connected to a number of mapping functions, where the exact mapping chosen depends on the client and rate. This allows both ODU0 mapping and mapping into a group of 155-Mbps timeslots, thereby using the bandwidth very efficiently. Two cross connects (XC) that work on the ODU0 and ODU1 level sit in the center of the TPO124. This enables add/drop/pass-through/protection of individual ODU0/ODU1 containers.

Figure 3. Block Diagram of TPO124: Any Client to 10G OTN Mapper



**Table 1** compares the features of the two described solutions. The TPO124 is based around the targeted feature set required to build an OTU1/OTU2 muxponder. Combined with the fact that the TPO124 is based on an Altera® low-cost 40-nm Arria® II GX FPGA, the TPO124 is designed to be the lowest cost and power solution for this application. If additional features are required (such as a packetized interface or a special enhanced FEC), the TPO124 can be adapted to include these new features. TPACK has many years of experience in making such customer-specific adaptations, with a typical turn-around time of three to six months. In contrast to the TPO124, the standard cell-based ASSP solution is a superset of features that may or may not be required for a given application. Most ASSPs developed today are in the best case based on 65-nm ASIC technology. Any changes, such as if a standard is changed or if a customer requests a specific feature, require a costly respin of the ASIC, which may take 9 to 12 months to complete.

*Table 1. Feature Comparison of PLD-Based OTN Mapping ASSPs and Custom ASIC-Based OTN Mapping ASSPs*

Features	TPACK TPO124	Standard Cell ASIC-Based ASSP
Process technology	40 nm	65 nm (or older/larger)
Client ports	8	8
Line ports	1	1
FEC	Standard FEC plus flexibility to add any proprietary enhanced FEC	Standard FEC plus two fixed specific enhanced FECs
ODU0 support	Yes	No
Ethernet-over-SONET and VCAT support	Not in standard device but can be added as needed	Yes
Packetized Ethernet and OTN support	Not in standard device but can be added as needed	Yes
Power	7–10 W	Unspecified
Relative ASP	2X–4X	4X–5X
Future updates as standards get ratified	ITU-T standardized OC-3/OC-12 mapping to OTN	None without ASIC respin

The flexibility and advanced feature set in the TPACK TPO124 is due in large part to TPACK's SOFTSILICON approach to implementing its device solutions, which is a complete standard-featured turnkey image for download in an Arria II GX FPGA.

## Implementation Options for OTN Mapping Solutions

The traditional path for solutions vendors addressing opportunities of this magnitude is to pursue a standard cell ASIC-based ASSP. These ASSPs are typically based on older semiconductor processes, including 65 nm and larger geometries, which are fairly stable and provide a very low cost structure. However, more innovative approaches are available through other solutions such as programmable logic devices (PLDs). Manufacturers of these devices are able to utilize more advanced semiconductor processes owing to their ability to amortize their costs across a broader range of customers. For example, Altera's most advanced FPGAs and HardCopy® ASICs are available at the 40-nm process node. They enable manufacturers to deliver the associated cost and power benefits to their customers earlier than those who rely on older process nodes.

**Table 2** shows the percentage of ASIC design starts at each process node, starting with the year 2002 and projected through to 2011. As shown by the red-outlined boxes, the 130-nm node was the number-one node for ASIC design starts beginning in 2003 and is projected to continue to be so until 2010. In comparison, the leading FPGA process (represented by Altera in the blue boxes) diverges in 2005, and continues to the present where Altera's 40-nm FPGAs are four process nodes ahead of the most popular ASIC design start process, and at least one process node ahead of nearly all ASIC design starts. This means that ASSP solutions based on Altera 40-nm FPGAs and HardCopy devices reap all of the associated cost, performance, and power-reduction advantages of being on a more advanced process node compared to ASSPs based on standard cell ASIC approaches.

Table 2. Percentage of ASIC Design Starts (red) at Each Process Node Compared with Leading FPGA Process Technology from Altera (blue)(1)(2)

Process Node	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
0.022 $\mu\text{m}$	0	0	0	0	0	0	0	0	0	0
0.032 $\mu\text{m}$	0	0	0	0	0	0	0	0	0	1
0.040 $\mu\text{m}$	0	0	0	0	0	0	0	2	4	5
0.045 $\mu\text{m}$	0	0	0	0	0	1	1	2	4	5
0.065 $\mu\text{m}$	0	0	0	1	2	5	9	10	13	15
0.09 $\mu\text{m}$	0	1	8	13	18	23	24	25	24	24
0.13 $\mu\text{m}$	18	37	42	29	29	27	25	25	24	23
0.18 $\mu\text{m}$	38	27	23	20	17	14	12	10	10	8
0.25 $\mu\text{m}$	16	15	12	12	11	9	9	8	7	5
0.35 $\mu\text{m}$	21	16	12	12	11	10	9	8	6	5
0.5 $\mu\text{m}$	5	4	3	7	7	6	6	5	4	4
L> 0.5 $\mu\text{m}$	1	1	0	6	6	6	5	5	5	5
Total	100	100	100	100	100	100	100	100	100	100


**Notes:**

- (1) Source: Gartner, published March 2009
- (2) Columns 2002 to 2004 show design starts in the Americas, columns 2005 to 2011 show design starts worldwide.

Despite the higher cost-per-logic of PLDs compared to standard cell ASICs, implementing standard products based on PLDs can result in lower-priced solutions. The cost savings of the PLD-based approach mainly derives from two sources. The process node advantage (illustrated in Table 2) enables PLD-based solutions to achieve the lowest die area per functionality. In addition, PLD-based solutions can be customized to include only the functions required by a certain application, whereas a standard cell ASIC-based ASSP is a superset of the features needed by customers in the solution space. The combination of these two factors enables PLD-based solutions to support low price points that are competitive or better than those of standard cell ASIC-based ASSPs. Further, the silicon cost associated with more advanced nodes has greater potential for the greatest decrease over time compared to older nodes, resulting in lower manufacturing costs in the long term for PLD-based products.

Going forward, the development cost for advanced process nodes continues to grow, as shown in Figure 4. These rising costs place this class of semiconductor device development increasingly out of reach for many component and equipment vendors. The rising costs also raise the bar for the return on the development investment, which increasingly must be derived from fewer customers as the components become more specialized. These dynamics in turn force the pricing of standard cell ASIC-based ASSPs products upward.

Figure 4. Cost of Development at Various Semiconductor Process Nodes



Another aspect favoring PLD-based solutions is time to market. As an example, TPO124 offers support for the newly standardized ODU0, which is not supported by any standard cell ASIC-based ASSPs at present. This demonstrates that new features can be brought to market much faster with an FPGA-based approach than with a standard cell ASIC-based approach.

TPACK uses its SOFTSILICON method, where a standard chip product (similar to an ASSP) is offered, to leverage the FPGA-based benefits. The product feature set is pre-defined with some configuration options available, and the product is updated regularly with new features to keep up with market developments. In the case of the TPO124 OTN mapper product, TPACK leverages the industry's lowest cost 40-nm FPGAs, the Arria II GX family, to help achieve price points that are competitive or lower than those offered by ASSP vendors who rely on standard cell ASIC-based solutions.

## Summary

Arria II GX FPGAs have helped Altera's partner TPACK develop and deliver an OTN mapper chip solution that can compete with a standard cell ASIC-based solution on cost, power, and features. The memory-to-logic density benefits of Arria II GX FPGAs match the needs of these applications, making the Arria II GX family a very cost-effective option. Additionally, the 3.75-Gbps-transceiver channels in Arria II GX FPGAs are well aligned with the key client port rates such as GbE, OC-12/OC-3, and Fibre Channel. When a product such as the TPO124 is based on low-cost 40-nm FPGA or HardCopy ASIC technology, it means that innovative developers can offer flexible semiconductor products with lower cost and lower power than alternative solutions based on standard cell-based ASICs.

## Further Information

- Arria II GX FPGA Family:  
[www.altera.com/products/devices/arria-fpgas/arria-ii-gx/aiigx-index.jsp](http://www.altera.com/products/devices/arria-fpgas/arria-ii-gx/aiigx-index.jsp)
- TPACK OTN Mappers:  
[www.tpack.com/solutions-and-products/smartpack-products/p-ocket-otn-mappers.html](http://www.tpack.com/solutions-and-products/smartpack-products/p-ocket-otn-mappers.html)
- Understanding SOFTSILICON:  
[www.tpack.com/products/softsilicon.html](http://www.tpack.com/products/softsilicon.html)

## Acknowledgments

- Thomas P. Rasmussen, Vice President of Product Line Management, TPACK
- Martin S. Won, Senior Member of Technical Staff, Altera Corporation



101 Innovation Drive  
San Jose, CA 95134  
[www.altera.com](http://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.