White Paper

Enhancing High-Speed Telecommunications Networks with FEC

As the demand for high-bandwidth telecommunications channels increases, service providers and equipment manufacturers must deliver more bandwidth for less cost. High-speed fiber optic links between towns, cities, and countries are expected to deliver bit rates of 10 Gbps or more, and new techniques are needed to meet these expectations. For example, dense wavelength division multiplexing (DWDM), which combines different wavelengths of light onto a single fiber, allows 40 Gbps capacity to be transmitted on a single fiber.

This increased bandwidth is not cheap. The optical components used in these systems are still a significant portion of the total system cost—even when shared among thousands of end users. Additionally, the customers and service providers expect a high-performance, error-free system that is reliable and always available.

The demand for increased bandwidth, high quality, and low cost has caused telecommunications standards bodies to investigate forward error correction (FEC) techniques—which are currently used in wireless and mass storage applications—for use in telecommunications systems. With FEC, systems can achieve higher performance by detecting and correcting errors on the link. Two future ITU specifications recommend FEC in transmission systems, ITU-T G.709 and ITU-T G.975. ITU-T G.709 defines the network node interface for the optical transport network operating at 2.5, 10 and 40 Gbps, and describes a "wrapper" approach that incorporates a simple framing structure and an FEC section. Specifically intended for use in optical fiber submarine cable systems, ITU-T G.975 recommends FEC in systems operating at 2.5 Gbps and higher.

Both specifications use a Reed-Solomon encoder/decoder for FEC. This white paper describes the reasons for this choice and describes a practical implementation of FEC for ITU-T G.709 systems using Altera® programmable logic devices (PLDs).

Advantages of FEC

The quality of a fiber optic link is determined by a variety of parameters. The span of a link is typically determined by the optical power budget, which is the difference in power between what the optical transmitter can produce and what the optical receiver can detect. Within the link, the attenuation from every kilometer of fiber and every connector or coupler adds together to consume this budget. To create links with large spans, system designers can choose to use amplifiers or repeaters. A system with a larger optical power budget preserves the link's bit error rate (BER) and has less need for amplifiers and repeaters. Designers typically achieve this increased optical power budget by using higher quality, higher cost optical components.

FEC effectively adds a significant gain to the link's optical power budget, while keeping the BER as low as possible. Because FEC systems must detect an error before correcting it, using FEC across a link lets the designer measure performance and identify link degradation early.

FEC provides the following benefits for fiber optic communications links:

- Improves performance of an existing link between two points
- Increases the maximum span of the link in systems without repeaters
- Increases the distance between repeaters in optically amplified systems or relaxes the specifications of the optical components or fiber
- Improves the overall quality of the link by diagnosing degradation and link problems earlier

The biggest disadvantage of using FEC is that the inserted check symbols consume bandwidth within the communications channel. A system using FEC typically requires a slightly higher bit rate to carry this additional correction data. The system designer's challenge is to select an FEC method that is cost effective, scalable, and able to meet the required high-speed bit rates.

Reed-Solomon Codes

Reed-Solomon codes are described as *(N,K)* where *N* is the total number of symbols per codeword, *K* is the number of information symbols, and *R* is the number of check symbols *(N-K)*. For example the Reed-Solomon codes used in ITU-T G.709 and G.975 are both (255,239) so will consist of a 239 information symbols and 16 check symbols.

Reed-Solomon codes use Galois field arithmetic constructed with a user-defined polynomial. The size of the Galois field is determined by the symbol width. In this example, the symbol width is 8 bits and the polynomial is defined as:

$x^8 + x^4 + x^3 + x^2 + 1$

Reed-Solomon codes treat errors on a symbol basis, therefore, a symbol that has all bits in error is as easy to detect and correct as a symbol that has a single bit in error. Reed-Solomon codes let the system detect and correct one error symbol for every two check symbols. For this specific example, up to 8 error symbols can be corrected per codeword, which is very effective for systems that operate with error rates much lower than 8/255.

Interleaving data from different codewords improves the efficiency of Reed-Solomon codes because the effect of burst errors is shared across many codewords. Therefore, interleaving codewords provides a coding gain for burst errors that is equivalent to the depth of interleaved codewords. For example, if 16 discrete encoders and decoders were used at either end of a channel, the maximum length of consecutive bit errors that could be detected and corrected would be 1,024. Both ITU-T G.709 and G.975 specify interleaving as part of the transport frame, which is useful for increasing the error correction and has the additional benefit of making the hardware implementation easier.

Advantages of Using Reed-Solomon Codes for FEC

Reed-Solomon is a block code, which means that whole blocks of symbols are encoded together, rather than single bytes or words. This technique has two key advantages for telecommunications:

- Most data transmitted across these high-bandwidth links is encapsulated in SONET or SDH frames. Therefore, the data is byte oriented and easily synchronized on 8-bit symbol boundaries.
- A block-based coding scheme makes a solution easily scalable. Block-based coding schemes can easily achieve a much higher throughput by employing parallel architectures that time-division multiplex many blocks of hardware together. This approach is used to implement the 10 Gbps FEC solution discussed on page 3.

Reed-Solomon is a very efficient coding scheme, measured in terms of number of errors corrected versus the algorithm complexity. For a relatively small amount of silicon, Reed-Solomon codes can correct an acceptable number of errors at relatively high transmission rates.

ITU-T G.709 Framing Structure

ITU-T G.709 specifies a framing structure that can handle data rates at 2.5, 10, and 40 Gbps. The highest framing level is referred to as the optical transport unit (OTU). The OTU*k* frame is identical in structure for all three data rates. *k* is defined as 1, 2, and 3 to denote 2.5, 10, and 40 Gbps, respectively. See Figure 1.

Figure 1. OTU2 Frame

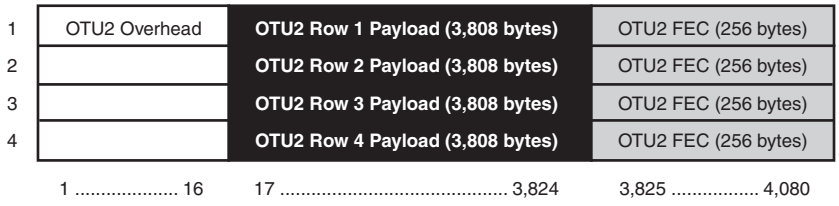

Overhead bytes are used for frame and multiframe alignment and for carrying communications channels. Reed-Solomon check symbols are added to the end of each row for the symbols in the data section.

ITU-T G.709 recommends an interleaved system in which each frame is divided into 16 interleaved sub-frames. This system makes encoding and decoding the check symbols easier and increases the system's error correction ability. Figure 2 shows the interleaving.

Figure 2. Interleaving Structure for 16 Sub-Frames with Individual FEC Sections

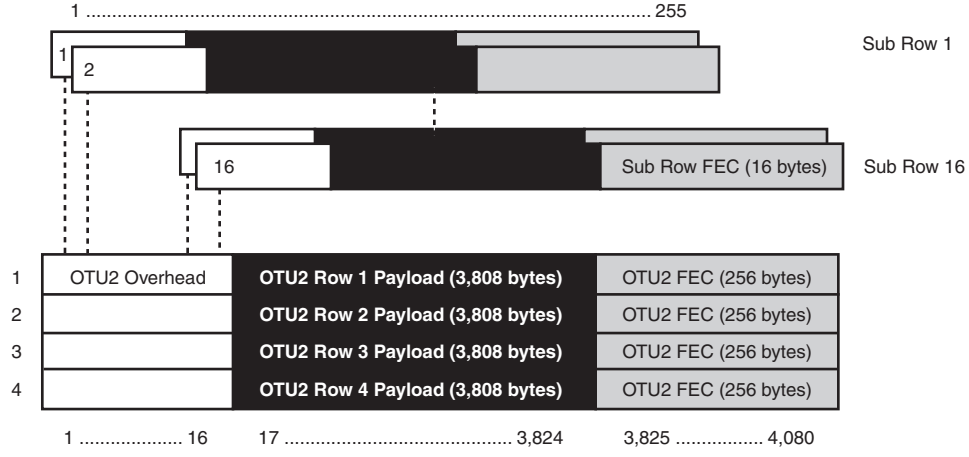

Designing a Reed-Solomon Decoder for OTU2 in Altera PLDs

ITU-T G.709 specifies interleaving the results from 16 Reed-Solomon encoders to build the OTU*k* frame. However, the maximum data rate that can be passed through a single encoder or decoder dictates the real number of encoders and decoders needed. For lower bit rates (2.5 Gbps) it is possible to use fewer encoders/decoders and store the encoded/decoded blocks in memory. APEX™ embedded system blocks (ESBs) are ideal for storing this data. Higher bit rates can require more than 16 discrete encoders and decoders with multiplexed outputs.

The bit rate of 10 Gbps systems converts to 9,953.2 Mbits/second. This data rate represents the incoming STM-64 frames or the aggregate rate of several lower order multiplexed paths. Adding the Reed-Solomon check symbols and simple framing bytes to these bytes as defined by ITU-T G.709—i.e., for every 238 data symbols one framing symbol and 16 check symbols are added--gives a bit rate of 10,619 Mbits/second.

If a single discrete Reed-Solomon decoder decodes at a rate of 800 Mbits/second for 8-bit symbols, the system theoretically needs 13.27 decoders. To make the clock generation and interleaving simpler, this number is rounded up to 16.

Building the Discrete Decoder

A decoder that meets the requirements of OTU2 must be as fast as possible and must be able to process a symbol on every clock edge. The following example system provides 10 Gbps and uses 16 discrete hardware decoders operating at an input clock rate of 83 Mhz and an output clock rate of 78 Mhz. This rate adaptation is handled by APEX ESBs, which are an integral part of the interleaving process shown in Figure 2.

The design uses the Reed-Solomon Compiler MegaCore® function continuous decoder variant. Using the MegaCore function shortens this stage of the design process from months to minutes and guarantees the best possible silicon efficiency and performance when targeting Altera devices. You can parameterize the MegaCore function using the MegaWizard® Plug-In interface to meet the requirements of ITU-T G.709. See Figure 3.

Figure 3. Setting Parameters to Meet ITU-T G.709 & G.975

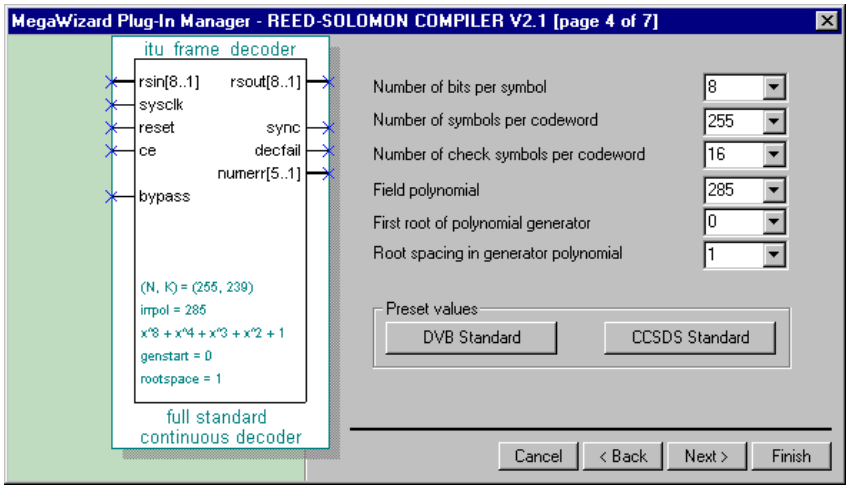

The wizard outputs a VHDL, Verilog HDL, or the Altera Hardware Description Language (AHDL) block that can be integrated with the rest of the design. This design example instantiates this block 16 times and has control and first-in first-out (FIFO) circuitry so the 16 decoders process the interleaved data stream correctly.

Building the Design

After building the single discrete decoder, the next step it to determine the additional circuitry needed to glue the system together. Figure 4 shows the architecture of the decoder and the clock domains that are decoupled by the FIFO buffers. Although this system works for 10-Gbps applications, you can scale it to support 2.5-or 40-Gbps applications (in the latter case by using multiple devices).

Figure 4. Building the Decoder Design

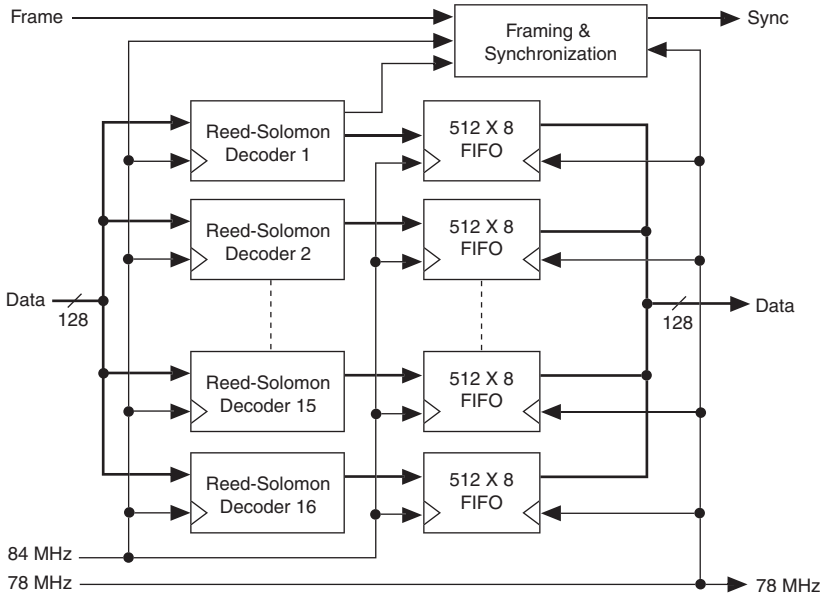

For the required parameters with the continuous decoder core, each discrete decoder requires 2,460 logic elements (LEs) and 7 ESBs. Additionally, FIFO buffers are needed at the output of the Reed-Solomon decoders to perform the rate adaptation. Each FIFO buffer requires 2 ESBs (512 x 8 FIFO). Therefore, 16 instances of the design require 40,000 LEs and 144 ESBs.

With these resource requirements, the design fits into an EP20K1500E device, leaving ample room for multiplexing functions, frame synchronization, and performance-monitoring counters. Table 1 gives utilization results for Reed-Solomon decoders implementing a FEC solution compliant to ITU-T G.709.

Table 1. Reed-Solomon Decoder Utilization

Data Rate Gbps	Resources Used		Device	Speed Grade	Number of Devices	Utilization	
	LEs	ESBs				LEs	ESBs
2.5	10,500	36	EP20K300E	-1	1	91%	50%
10	42,000	144	EP20K1500E	-1	1	81%	67%
40	168,000	576	EP20K1500E	-1	4	81%	67%

Conclusion

Although it is a new technology for these applications, FEC brings provides significant benefits to telecommunications networks. With PLDs, you can rapidly deploy telecommunications equipment featuring FEC, and still upgrade your system in the future to meet specification changes or updates. For more information on using the Reed-Solomon MegaCore function, refer to the ***Reed-Solomon Compiler MegaCore Function User Guide***.

101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
http://www.altera.com

Copyright © 2001 Altera Corporation. Altera, APEX, MegaCore, and MegaWizard are trademarks and/or service marks of Altera Corporation in the United States and other countries. Altera acknowledges the trademarks of other organizations for their respective products or services mentioned in this document. 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. All rights reserved.