

LogiCORE IP Viterbi Decoder v8.0

Product Guide

PG027 January 18, 2012

Table of Contents

Chapter 1: Overview

Standards Compliance	4
Feature Summary	4
Licensing	5
Performance Characteristics	5
Resource Utilization.....	8

Chapter 2: Core Interfaces

Port Descriptions.....	9
AXI4-Stream Protocol	10

Chapter 3: Customizing and Generating the Core

CORE Generator Parameters	17
Parameter Values in the XCO File	23
Output Generation	24

Chapter 4: Designing with the Core

Functional Description	25
Design Guidelines	33
Viterbi Decoder Non-features Summary.....	35

Chapter 5: Detailed Example Design

Demonstration Test Bench	36
--------------------------------	----

Appendix A: Migrating

Parameter Changes in the XCO File.....	38
--	----

Appendix B: Debugging

Appendix C: Additional Resources

Xilinx Resources	43
Solution Centers	43
References	43
Technical Support.....	43
Ordering Information	44
Revision History	44
Notice of Disclaimer	44

Introduction

The Viterbi Decoder is used in many Forward Error Correction (FEC) applications and in systems where data are transmitted and subject to errors before reception. The Viterbi Decoder is compatible with many common standards, such as DVB, 3GPP2, 3GPP LTE, IEEE 802.16, Hiperlan, and Intelsat IESS-308/309.

Features

- High-speed, compact Viterbi Decoder
- Fully synchronous design using a single clock
- Parameterizable constraint length from 7 to 9
- Parameterizable convolution codes
- Parameterizable traceback length
- Decoder rates from 1/2 to 1/7
- Very low latency option
- Minimal block RAM requirements; two block RAMs for a constraint length 7 decoder
- Serial architecture for small area
- Soft decision with parameterizable soft width
- Multichannel decoding
- Dual rate decoder
- Trellis mode
- Erasure for external puncturing
- BER monitor
- Normalization
- Best state option
- For use with Xilinx CORE Generator™ software and Xilinx System Generator for DSP v13.4
- Compatible encoder core available in the Xilinx CORE Generator™ software

LogiCORE™ IP Facts Table	
Core Specifics	
Supported Device Family ⁽¹⁾	Zynq™-7000, Artix™-7, Virtex® -7, Kintex™-7, Virtex-6, Spartan® -6
Supported User Interfaces	AXI4-Stream
Provided with Core	
Design Files	Netlist
Example Design	Not Provided
Test Bench	VHDL
Constraints File	Not Provided
Simulation Model	Verilog VHDL
Supported S/W Driver	N/A
Tested Design Tools	
Design Entry Tools	CORE Generator tool 13.4 System Generator for DSP 13.4
Simulation ⁽²⁾	Mentor Graphics ModelSim Cadence Incisive Enterprise Simulator (IES) Synopsys VCS and VCS MX ISim
Synthesis Tools	N/A
Support	
Provided by Xilinx @ www.xilinx.com/support	

1. For a complete listing of supported devices, see the [release notes](#) for this core.
2. For the supported versions of the tools, see the [ISE Design Suite 13: Release Notes Guide](#).

Overview

This core implements a Viterbi Decoder for decoding convolutionally encoded data. For detailed information on the design see [Chapter 4, Designing with the Core](#).

Standards Compliance

The Viterbi Decoder core adheres to the AMBA® AXI4-Stream standard [\[Ref 5\]](#).

Feature Summary

In modern communication systems, there is a requirement to transmit data and recover it, without error, in the presence of noise. This prevents having to retransmit the data, if there are errors, which would reduce the data rate in the system. One technique used is convolutional coding. A convolutional encoder [\[Ref 2\]](#) and Viterbi Decoder [\[Ref 3\]](#) are used together to provide the error correction. The convolutional encoder adds redundancy to the original data, and in the presence of noise the Viterbi Decoder uses maximum likelihood decoding to recover the data.

The convolutional encoder encodes the input data. A typical code rate for an encoder is 1/2, which signifies that for each input bit there are two output bits from the encoder. Similarly, for a code rate 1/3, each input bit has three output bits. Generator polynomials are used to encode each output bit from the convolutional encoder, thereby providing error protection for the input data. The encoder implementation consists of XOR gates and shift registers.

The Viterbi Decoder is configured to the same parameters as the encoder - code rate, constraint length, and the generator polynomials. The format of the input data to the Viterbi Decoder can be either hard or soft coding. A hard code is a binary value, whereas a soft code has a number of levels to reflect the strength, and hence confidence level, of the input data. This allows the Viterbi Decoder to know how strong '1' or '0' can be, which results in a better error protection. The output of the Viterbi Decoder is the original data that was input into the encoder.

The summary features for the Viterbi Decoder are as follows:

- Parameterizable decoder rates, constraint length, convolution codes and traceback lengths
- Choice of either parallel architecture for high data throughput, or serial for smaller area footprint
- Very low latency option
- Soft decision with parameterizable soft width

- Other architectural options such as multichannel decoding, dual rate decoder or trellis mode
- Erasure for external puncturing

Licensing

The Viterbi Decoder core is provided under the [SignOnce IP Site License](#) and can be generated using the Xilinx® CORE Generator™ v13.4. The CORE Generator software is shipped with ISE® Design Suite software v13.4.

To access the full functionality of the core, including simulation and FPGA bitstream generation, a full license must be obtained from Xilinx. For more information, visit the Viterbi Decoder [product page](#).

Performance Characteristics

It is important to set a maximum period constraint on the core clock input. The data in [Tables 1-1](#) show clock speeds that can be achieved when this is done. It might be possible to improve slightly on these values by trying different options for the place and route software. If necessary, performance can be increased by selecting a part with a faster speed grade.

Table 1-1: Viterbi Decoder Characterization Data⁽¹⁾

	Parallel	Serial	Multichannel 3 Channels
Xilinx Part	xc7vx330t	xc7vx330t	xc7vx330t
LUT/FF Pairs	2903	1771	2559
LUTs ^[4]	2525	1532	2312
FFs	1116	1915	3258
Block RAMs (36k) ^[5]	2	2	2
DSP Blocks	0	0	0
Max Clock Freq ^{[2][3]}	286/403	311/482	342/458
Mb/s	286/403	25/40	114/152 per channel

Notes:

1. Results shown for Viterbi Decoder with Constraint Length 7, Output Rate 1/2, Traceback 96, Soft Width 3, and best state on Virtex®-7 FPGA.
2. Area and maximum clock frequencies are provided as a guide. They can vary with new releases of the Xilinx implementation tools.
3. Maximum clock frequencies are shown in MHz for -1/-3 parts for Virtex-7 FPGAs. Clock frequency does not take jitter into account and should be de-rated by an amount appropriate to the clock source jitter specification.
4. LUT count includes route-thrus and can vary when the core is packed with other logic.
5. This is the total number of 36k block RAMs used when map was run. In reality, two 18k block RAM primitives can usually be packed together, giving an absolute minimum total block RAM usage of block RAMs (36k) + (block RAMs (18k) / 2) (rounded up).

Latency

The latency of the core depends on the traceback length and the constraint length. If the reduced latency option is selected, then the latency of the core is approximately halved and the latency is only 2 times the traceback length. The latencies given in the following sections are a count of the number of symbol inputs between `s_axis_data_tdata` and the decoded data result on the output `m_axis_data_tdata`. The actual latency depends on the parameters selected for the core and the true value of the latency for a given set of parameters can be found through simulation. The `tvalid` signal indicates when there is valid data on the output of the core.

Without Reduced Latency

For a parallel core, the latency is of the order

$$\text{Latency} \cong 4 * \text{traceback_length} + \text{constraint_length} + \text{output_rate}$$

For a serial core, the latency is of the order shown. Note that the latency is in terms of valid inputs in the serial case.

$$\text{Latency} \cong 4 * \text{traceback_length} + \text{constraint_length}$$

With Reduced Latency

Reduced latency is only available for the parallel core. The latency is given by

$$\text{Latency} \cong 2 * \text{traceback_length} + \text{constraint_length} + \text{output_rate}$$

Multichannel Latency

The multichannel core always uses the standard latency option. The latency in the multichannel case is given by

$$\text{Latency} \cong 4 * \text{traceback_length} * \text{channel_count} + (\text{constraint_length} * \text{channel_count}) + \text{output_rate}$$

This corresponds to the reduced latency single-channel case above with channel count set to 1.

Trellis Mode Latency

The latency of the Trellis Mode Decoder is as the equations above, but reduced by the `output_rate` as the branch metric costing unit is not present in the core.

BER Performance

BER performance curves were generated using a hardware-in-the-loop test framework, consisting of the Xilinx Convolution Encoder v8.0, an AWGN channel model, the Viterbi Decoder and logic to implement BPSK modulation, soft data mapping and data comparison and collection. Figure 1-1 shows the basic system dataflow.

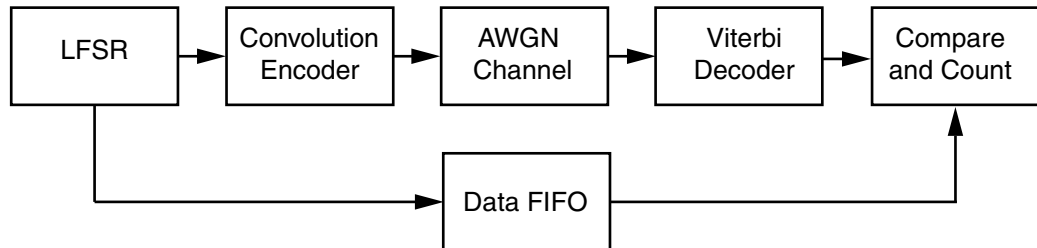


Figure 1-1: BER Performance Test System Data Flow Diagram

Figure 1-2 shows the BER performance of the core for constraint lengths 7 and 9 with soft width 4 and rate 1/2.

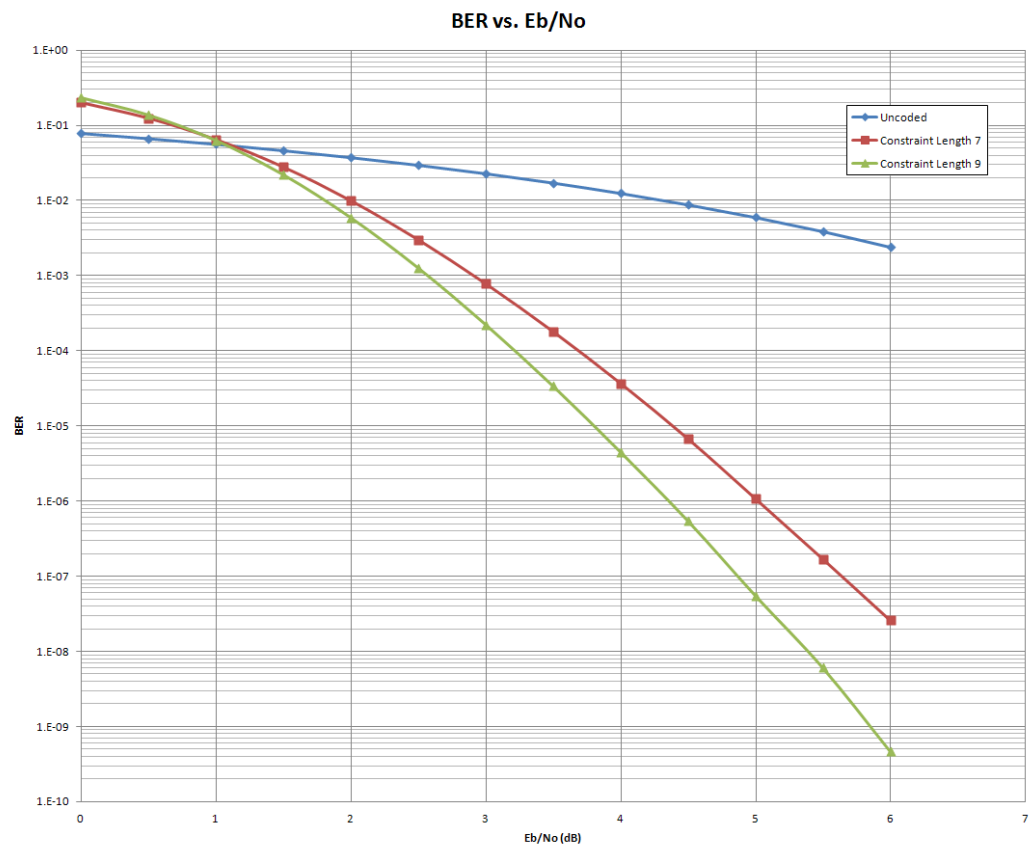


Figure 1-2: BER Performance for Rate 1/2 System

Resource Utilization

The area of the core increases with the constraint length and the soft width of the input data. Some example configurations are shown in [Performance Characteristics, page 5](#). The slice counts can be reduced slightly by selecting the option to map primary I/O registers into IOBs during placement. This option should certainly be selected if the core I/Os are to be connected directly onto a PCB via the FPGA package pins. This gives lower output clock-to-out times and predictable setup and hold times.

Block RAM Utilization

The block RAM requirements of the core depend on the constraint length and the traceback length. If the reduced latency option is selected, an extra block of RAM is required for the traceback addressing. Multichannel cores also require extra traceback as each channel requires its own traceback; thus the internal traceback length of the multichannel core is (channel count * traceback length). See [Table 1-2](#) for block RAM usage.

Table 1-2: Block RAM Requirements for the Viterbi Decoder with Standard Latency

Constraint Length	Block RAM
7	2
8	4
9	8

Core Interfaces

This chapter provides detailed descriptions for each interface.

Port Descriptions

A representative symbol of the Viterbi Decoder, with the signal names, is shown in [Figure 2-1](#) and [Figure 2-2](#) and described in [Table 2-1](#). Some of the pins are optional. These should be selected only if they are genuinely required, as their inclusion might result in an increase in the core size. Timing diagrams for the signals are shown in [Figures 2-3](#) to [2-6](#).

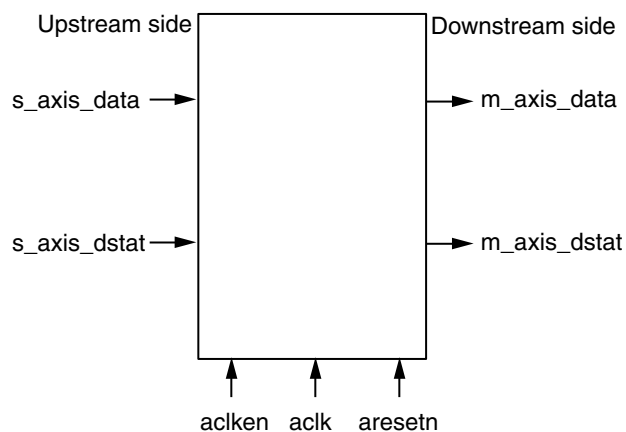


Figure 2-1: Core AXI Channels

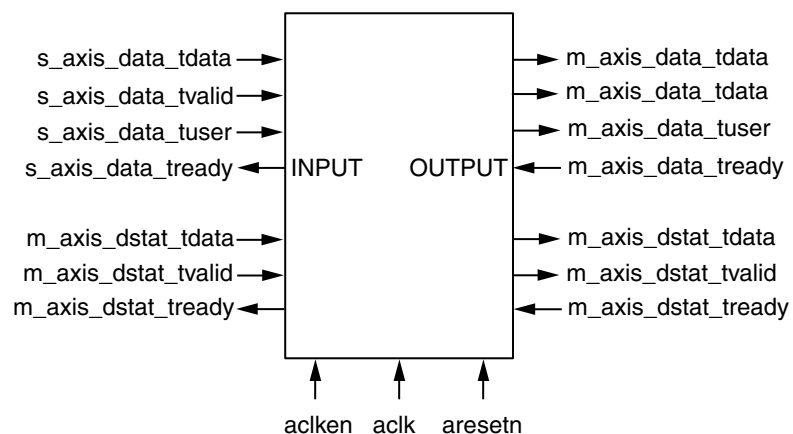


Figure 2-2: Core Schematic Symbol

Table 2-1: Signal Descriptions

Signal	Direction	Description
aclk	Input	Rising edge clock
aclken	Input	Active High clock enable (optional)
aresetn	Input	Active Low synchronous clear (overrides aclken)
s_axis_data_tdata	Input	Input data
s_axis_data_tvalid	Input	tvalid for S_AXIS_DATA channel. See AXI4-Stream Protocol .
s_axis_data_tready	Output	tready for S_AXIS_DATA. Indicates that the core is ready to accept data.
m_axis_data_tdata	Output	tdata for the output data channel, decoded output data.
m_axis_data_tvalid	Output	tvalid for M_AXIS_DATA channel.
m_axis_data_tready	Input	tready for M_AXIS_DATA channel. Do not enable optional tready pins or tie port high if downstream slave is always able to accept data from M_AXIS_DATA.
s_axis_dstat_tdata	Input	Input status data; for the Viterbi Decoder this is the BER count.
s_axis_dstat_tvalid	Input	tvalid for S_AXIS_DSTAT channel. See AXI4-Stream Protocol .
s_axis_dstat_tready	Output	tready for S_AXIS_DSTAT. Indicates that the core is ready to accept data. Always high, except after a reset if there is not a tready on the output.
m_axis_dstat_tdata	Output	tdata for the output DSTAT channel. Outputs the number of errors on the channel.
m_axis_dstat_tvalid	Output	tvalid for M_AXIS_DSTAT channel.
m_axis_dstat_tready	Input	tready for M_AXIS_DSTAT channel. Do not enable or tie high if downstream slave is always able to accept data from M_AXIS_DSTAT.

AXI4-Stream Protocol

The use of AXI4-Stream interfaces brings standardization and enhances interoperability of Xilinx® IP LogiCORE™ solutions. Other than general control signals such as aclk, aclken and aresetn, and event outputs, all inputs and outputs to the core are conveyed using AXI4-Stream channels. A channel consists of tvalid and tdata always, plus several optional ports and fields. In the Viterbi Decoder, the additional ports used are tuser and tready. Together, tvalid and tready perform a handshake to transfer a value, where the payload is tdata. The payload is indeterminate when tvalid is deasserted.

The Viterbi Decoder operates on the values contained in the S_AXIS_DATA channel tdata fields and outputs the results in the tdata fields of the M_AXIS_DATA channel. For further details on AXI4-Stream Interfaces see [\[Ref 4\]](#) and [\[Ref 5\]](#).

Basic Handshake

[Figure 2-3](#) shows the transfer of data in an AXI4-Stream channel. tvalid is driven by the source (master) side of the channel and tready is driven by the receiver (slave). tvalid

indicates that the value in the payload fields (`tdata` and `tuser`) is valid. `tready` indicates that the slave is ready to receive data. When both `tvalid` and `tready` are true in a cycle, a transfer occurs. The master and slave set `tvalid` and `tready` respectively for the next transfer appropriately.

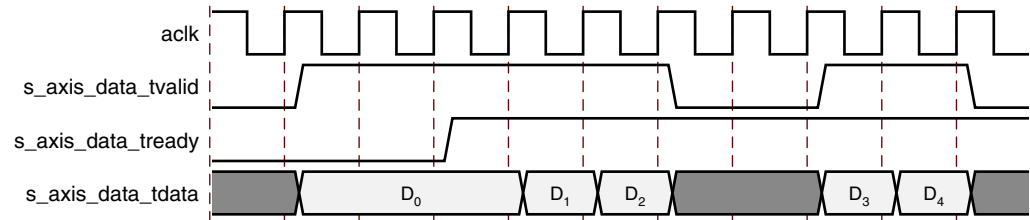


Figure 2-3: Data Transfer in an AXI4-Stream Channel

The full flow control of AXI4-Stream aids system design because the flow of data is self-regulating. Data loss is prevented by the presence of back pressure (`tready`), so that data is only propagated when the downstream datapath is ready to process it.

For the main output channel, `M_AXIS_DATA`, if the output is prevented from off-loading data because `m_axis_data_tready` is low then data accumulates in the core. When the core's internal buffers are full the core stops further operations. When the internal buffers fill, the `tready` (`s_axis_data_tready`) is deasserted to prevent further input. This is the normal action of back pressure.

For the status output channel, `M_AXIS_DSTAT`, if `m_axis_dstat_tready` is held low, the core does not overwrite the internal BER value until the current value within the core has been read because there is no internal buffering on the `DSTAT` channel. The core holds `m_axis_dstat_tdata` static after `m_axis_dstat_tvalid` goes high until `tready` is asserted, see Figure 2-4.

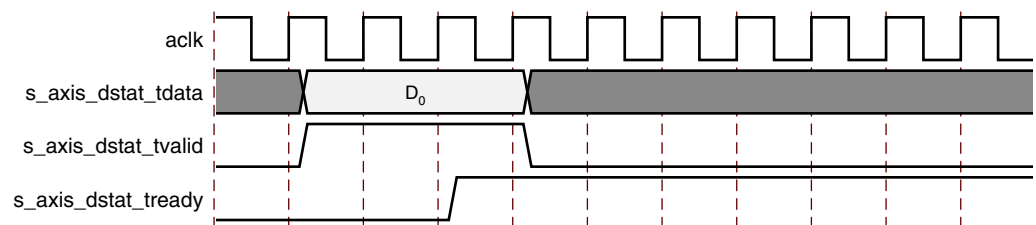


Figure 2-4: DSTAT Handshaking

aclken

The clock enable input (`aclken`) is an optional pin. When `aclken` is deasserted (low), all the other synchronous inputs are ignored, except `aresetn`, and the core remains in its current state. This pin should be used only if it is genuinely required because it has a high fanout within the core and can result in lower performance. `aclken` is a *true* clock enable and causes the entire core to freeze state when it is low.

An example of `aclken` operation is shown in Figure 2-5. In this case, the core ignores symbol D_4 as input to the block, and the current `m_axis_data_tdata` value remains unchanged.

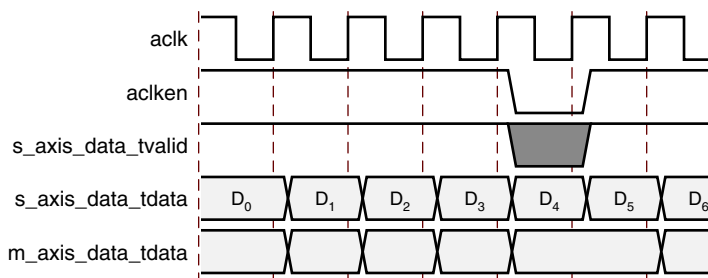
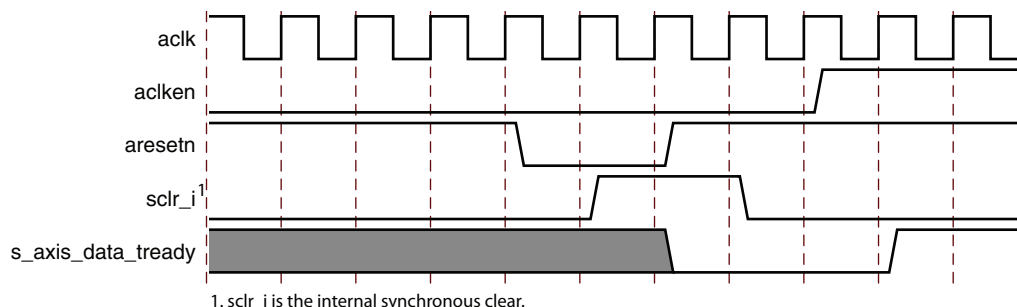


Figure 2-5: Clock Enable Timing

aresetn

The synchronous reset (`aresetn`) input can be used to re-initialize the core at any time, regardless of the state of `aclken`. `aresetn` needs to be asserted low for at least two clock cycles to initialize the circuit. The core becomes ready for normal operation two cycles after `aresetn` goes high, if `aclken` is asserted. Note that the block RAM is not cleared with the `aresetn` signal and there is a block of previously decoded data output from the traceback prior to correct decoding resuming. The `tvalid` signal on `m_axis_data` is only asserted when valid data is available on `m_axis_data_tdata`. The timing for the `aresetn` input is shown in Figure 2-6.



1. sclr_i is the internal synchronous clear.

Figure 2-6: Synchronous Reset Timing

S_AXIS_DATA Channel

s_axis_data_tdata

Data to be processed is passed into the core on this port. To ease interoperability with byte-oriented buses, `tdata` is padded with zeros because the Viterbi Decoder bus width can vary, depending on the type of Viterbi and the output rate. The padding bits are ignored by the core and do not result in additional resource use. The structure is shown in Figure 2-7 for an input rate 2 decoder. The `tdata` carries the data to be decoded. The encoded bits on each of the `data_in` inputs can be hard coded (bus width 1) or soft coded (bus width 3 to 5). The width of the input for a non-trellis decoder is always 8 times the output rate. If the Dual Decoder is selected, the number of `data_in` inputs is equal to the maximum output rate.

The input format for the trellis decoder is shown in Figure 2-8. The width of the trellis mode inputs can range from 4 to 6 corresponding to a data width of 3 to 5. The trellis mode

inputs are the outputs from an external costing of the data. There are always four inputs to the decoder, and the decoder always functions as a rate 1/2 decoder when trellis mode is selected. The width for the trellis decoder is always 40 bits with the sector input residing in the top byte. The SECTOR bus has width 4. The SECTOR input is delayed by the decoder delay and output to the SECTOR bus in `m_axis_data_tdata`. See [Trellis Mode Decoder, page 29](#). The following buses are available on `s_axis_data_tdata` for all decoders except the Trellis Mode Decoder:

DATA_IN0 input data which can be 1 bit for hard decoding or 3 to 5 bits wide for soft coding
 DATA_IN1 input data which can be 1 bit for hard decoding or 3 to 5 bits wide for soft coding
 DATA_IN2 input data which can be 1 bit for hard decoding or 3 to 5 bits wide for soft coding
 up to the output rate

The following buses are available on `s_axis_data_tdata` for the Trellis Mode Decoder:

TCM00, TCM01... TCM11 Input trellis data with width of 4 to 6 bits
 SECTOR Input sector of width 4 bits

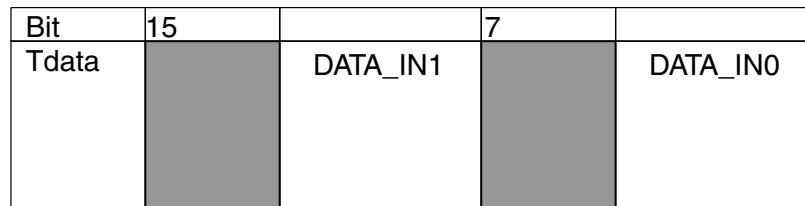


Figure 2-7: Input tdata for Standard Output Rate 2 Decoder with Soft Input

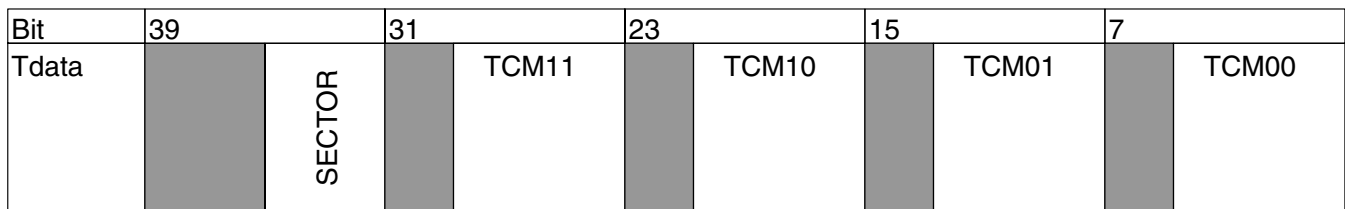


Figure 2-8: Input tdata for Trellis Mode Decoder

s_axis_data_tuser

This port is only present if the core is punctured, or is a Dual Decoder or the block valid signal is used with the core.

ERASE This erase input bus is only required where data on the channel has been punctured. The inputs are used to indicate the presence of a null-symbol on the corresponding `data_in` buses. `ERASE(0)` corresponds to `DATA_IN0`, `ERASE(1)` corresponds to `DATA_IN1`, ... If an erase pin is high, the data on the corresponding `data_in` bus is treated as a null-symbol internally to the decoder. The width of the erase bus is equal to the output rate of the decoder with a maximum value of 7.

SEL	Optional, controlled by the Dual Decoder option. This is used to select the correct set of convolutional codes for the decoding of the input data symbols in the Dual Decoder case. When SEL is low, the input data is decoded using the first set of convolutional codes. When it is high, the second set of convolutional codes is applied. See Figure 4-6 .
BLOCK_IN	Optional, controlled by the BLOCK VALID option. Marker signal used to tie output to input. The BLOCK_IN pin is delayed by the latency of the decoder and output as BLOCK_OUT on the M_AXIS_DATA_TUSER bus. The BLOCK_OUT pin shows the decoded data corresponding to the original BLOCK_IN set of data points.

The width of `s_axis_data_tuser` is always a multiple of 8 bits and is determined by the presence or absence of the three signals above.

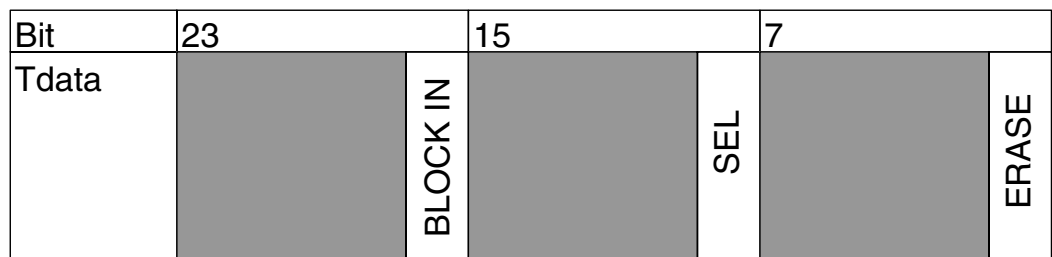


Figure 2-9: Input TUSER

M_AXIS_DATA Channel

m_axis_data_tdata

This output bus is composed of multiple fields:

DATA	Decoded output data always 1 bit.
SECTOR	This output is present if the decoder is a Trellis Mode Decoder and is always 4 bits. The output SECTOR is a delayed version of the input SECTOR bus. Both buses have a fixed width of 4 bits. The delay equals the delay through the Trellis Mode Decoder.

The width of the `m_axis_data_tdata` is 8 bits if the core is not a trellis decoder and 16 bits otherwise. See [Figure 2-10](#).

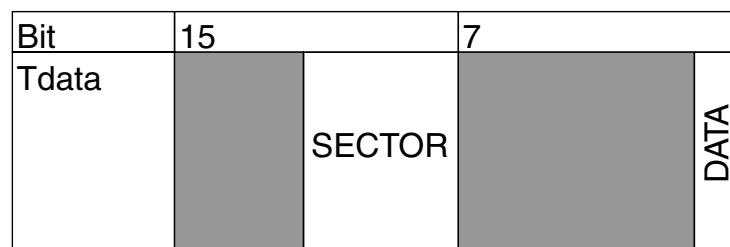


Figure 2-10: tdata Output

m_axis_data_tuser

This port is only present if the core is a Dual Decoder or has a normalization signal or block valid is present.

SEL	This signal is a delayed version of the SEL signal. The delay equals the delay through the Dual decoder.
BLOCK_OUT	This signal is a delayed version of the BLOCK_IN signal. The BLOCK_OUT signal shows the decoded data corresponding to the original BLOCK_IN set of data points. The delay equals the delay through the decoder.
NORM	The NORM output indicates when normalization has occurred within the core. It gives an immediate indication of the rate of errors in the channel. See Normalization, page 32 for additional details.

The width of m_axis_data_tuser is always a multiple of 8 bits and is determined by the presence or absence of the three signals above.

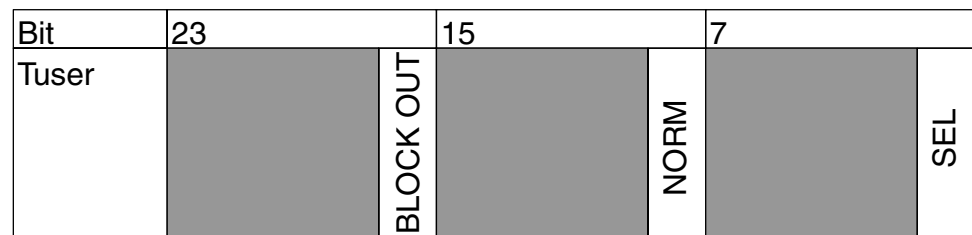


Figure 2-11: TUSER Output

S_AXIS_DSTAT Channel

s_axis_dstat_tdata

BER_RANGE This is the number of symbols over which errors are counted in the BER block, see [BER, page 32](#).

The width of s_axis_dstat_tdata is always 16 bits.

M_AXIS_DSTAT Channel

m_axis_dstat_tdata

BER The Bit Error Rate (BER) bus output (fixed width 16) gives a measurement of the channel bit error rate by counting the difference between the re-encoded DATA_OUT and the delayed DATA_IN to the decoder. For a full description of BER, see [BER, page 32](#). Trellis mode does not support a BER output.

Bit	15	7
Tdata	BER	

Figure 2-12: DSTAT Output

Customizing and Generating the Core

This chapter includes information on using Xilinx tools to customize and generate the core.

CORE Generator Parameters

Figure 3-1 shows the main CORE Generator™ Viterbi Decoder screen. To generate a core, click Generate.

Screen 1 (Viterbi Type)

Figure 3-1 shows the Viterbi Decoder Screen 1. The parameter descriptions for this screen follow.

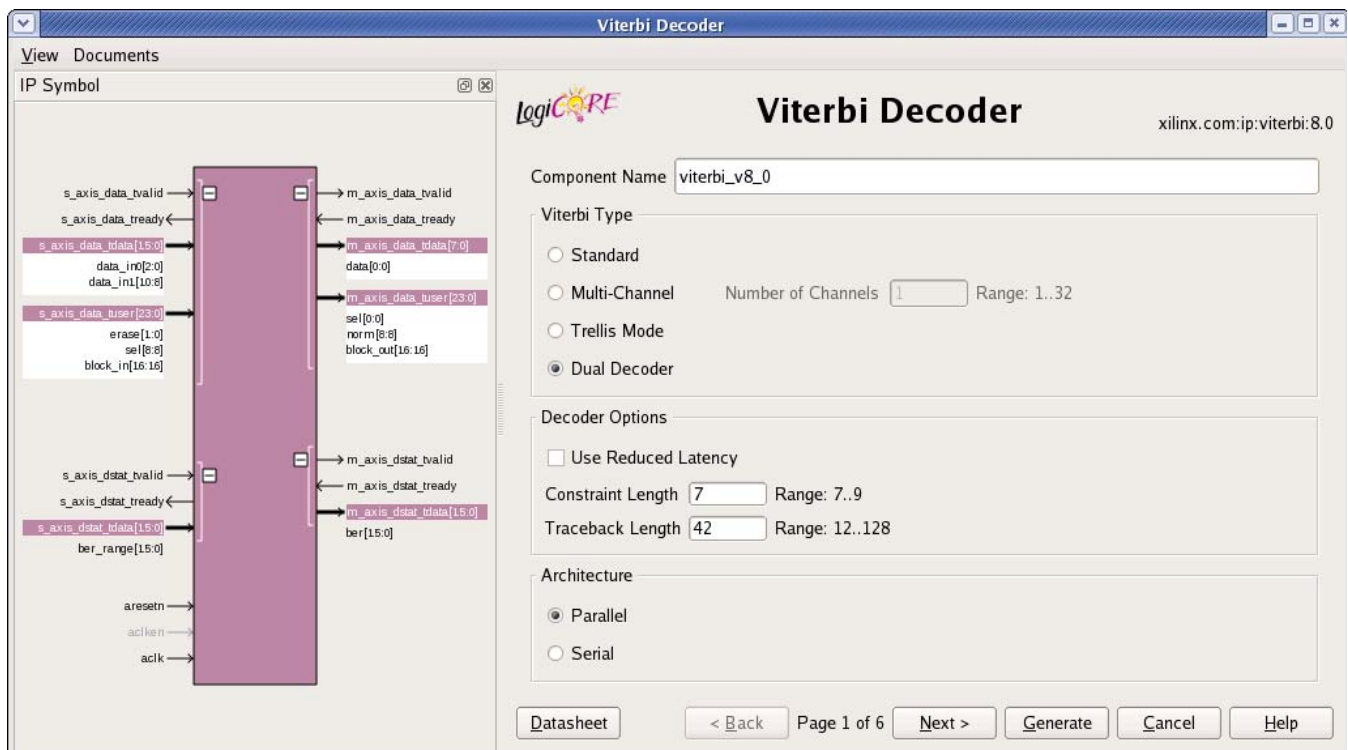


Figure 3-1: Viterbi Decoder Screen 1 (Viterbi Type)

Component Name

The component name is used as the base name of the output files generated for the core. Names must begin with a letter and must be composed of the following characters: a to z, 0 to 9, and “_”.

Viterbi Type

There are four different types of Viterbi Decoders which can be selected:

- **Standard:** This type is the basic Viterbi Decoder.
- **Multichannel:** This type allows many interlaced channels of data to be decoded using a single Viterbi Decoder. The number of channels to be decoded can be any value between 2 and 32. See [Multichannel Decoder, page 28](#).
- **Trellis Mode:** This type is a Trellis Mode Decoder using the costed data and SECTOR inputs. See [Trellis Mode Decoder, page 29](#).
- **Dual Decoder:** The type is the Dual Decoder that can be operated in dual mode with two sets of convolutional codes. The SEL pin is present on the slave and master TUSER buses when the decoder operates in this mode. See [Dual Rate Decoder, page 30](#).

Decoder Options

- **Constraint Length:** This is the length of the constraint register in the encoder plus 1. This value can be any integer in the range 7 to 9 inclusive.
- **Traceback Length:** This is the length of the survivor or training sequence in the traceback through the Viterbi trellis. Optimal length for the traceback is considered to be at least 6 times the constraint length for non-punctured data. For the multichannel Viterbi, the traceback length is the length of the traceback for each channel in the decoder. For the reduced latency option, the traceback length must always be divisible by 6. For punctured data, the length should be at least 12 times the constraint length. Increasing traceback length might increase RAM requirements. See [Block RAM Utilization, page 8](#) for more details. Increasing traceback length also increases the latency of the core. See [Latency, page 6](#) for additional details.
- **Use Reduced Latency:** This option reduces the latency on the core by approximately half. The reduced latency option has a slight speed penalty and is not available with the multichannel Viterbi. See [Latency, page 6](#) section for additional details.

Architecture

- **Parallel:** Large but fast Viterbi Decoder.
- **Serial:** Small but serial processing of the input data (see [Serial Decoder](#)).

Screen 2 (Architecture and Data Format)

See Figure 3-2. The parameter descriptions for this screen follow.



Figure 3-2: Viterbi Decoder Screen 2 (Best State, Puncturing and Data Format)

Best State

The best state option starts the traceback of the core from the optimal state.

- **Use Best State:** The best state selection gives improved BER performance for highly punctured data.
- **Best State Width:** The best state selects the best state from the costs for each state. Most of the lower bits in the cost are redundant in the cost comparison, and the best state area requirements can be reduced by selecting a smaller width than the full ACS width. If the width is set to 6, then the full cost is used in the best state selection for a soft width of 3. If the width is set to 3, then the lower three bits are ignored in the best state calculations.

Puncturing Options

- **None:** This indicates there is no external puncturing on the core.
- **External (Erased Symbols):** This indicates the presence of the erased bus ERASE on the TUSER input and allows the core to be de-punctured externally prior to decoding. ERASE(0) high indicates that the sample on DATA_IN0 is a null symbol; ERASE(1) high indicates that the data on DATA_IN1 is a null symbol;... The size of the erase bus is equal to the output rate of the decoder, or the maximum output rate if the Dual Decoder is selected.

Coding

There are two types of coding available: Soft Coding and Hard Coding.

- **Soft Coding:** Uses the Euclidean metric to cost the incoming data against the branches of the Viterbi trellis.
- **Hard Coding:** Uses the Hamming difference between the input data bits and the branches of the Viterbi trellis. Hard coding is only available for the standard parallel core.

Soft Width

The input width of soft-coded data is in the range 3 to 5. Larger widths require more logic. If the core is implemented in serial mode, larger soft widths also increase the serial processing time. See [Table 4-3](#) for the minimum number of clock cycles required for a rate 1/2 decoder in the serial case.

Data Format

There are two data formats available for Soft Coding: Signed Magnitude and Offset Binary. [Table 4-1](#) shows the required format for the data for the case of soft width 3 for each of the data types. Soft width 4 and 5 follow a similar format.

Screen 3 (Output Rate and Convolution Code)

See [Figure 3-3](#). The parameter descriptions for this screen follow.

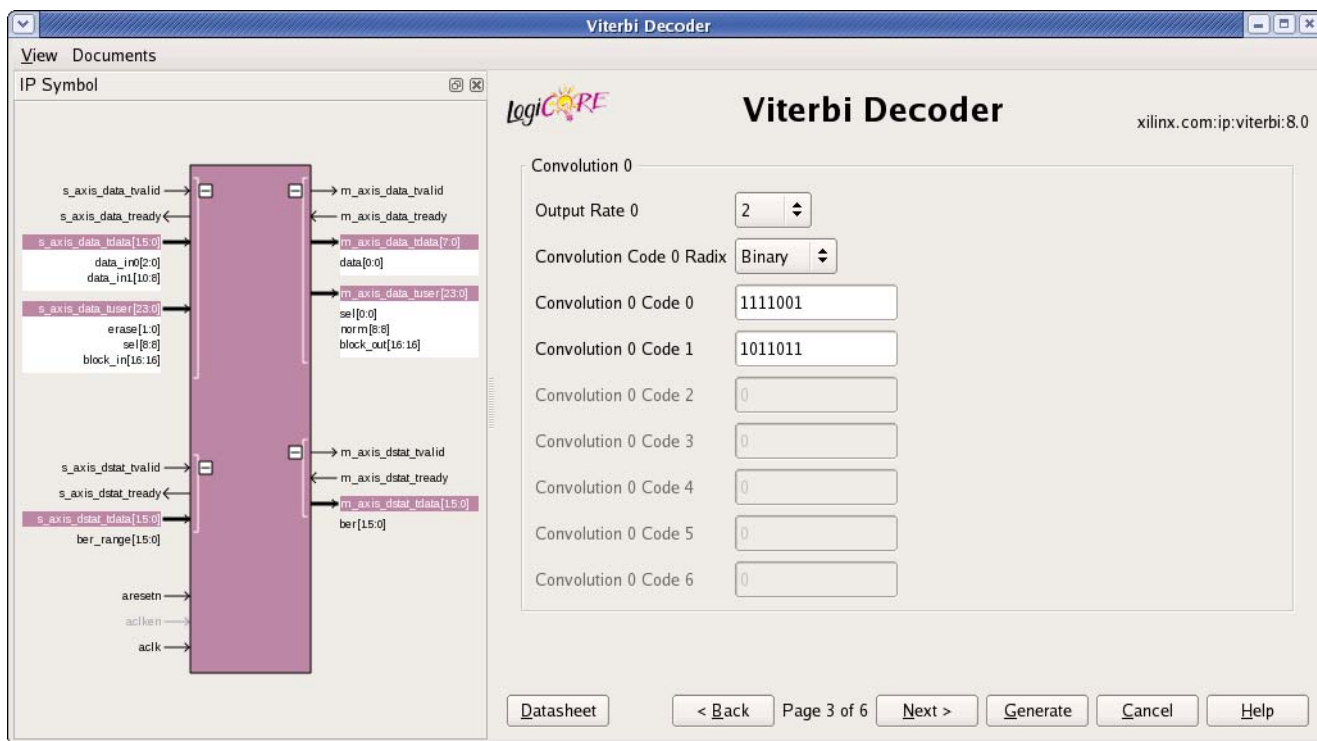


Figure 3-3: Viterbi Decoder Screen 3 (Output Rate and Convolution Codes)

Convolution Code0 Radix

The convolutional codes can be input and viewed in binary, octal, or decimal.

Output Rate0

Output Rate is the symbol output rate at the Encoder. Output Rate0 can be any value from 2 to 7. Output Rate0 is the output rate used if the decoder is non-dual. If the decoder is dual, then Output Rate0 is the first output rate and the rate used by the decoder when the SEL input is low.

Convolution0 Codes

These codes are the convolutional codes used in the encoder. The codes can be entered (and viewed) in binary, octal, and decimal. If the decoder is dual, then Convolution0 Codes are the codes applied in the decoder when the SEL input is low.

If the Dual Decoder type is selected, then a second output rate and convolution code selection screen is shown:

Output Rate1

Output Rate1 can be any value from 2 to 7. This is the second output rate used if the decoder is dual. The incoming data is decoded at this rate when the SEL input is high. Output Rate1 is not used for the non-Dual Decoder and the screen is only available if Dual Decoder is selected.

Convolution Code1 Radix

The convolutional codes can be input and viewed in binary, octal, or decimal.

Convolution1 Codes

The convolutional codes are used in the decoder when the decoder is dual and the SEL input is high. The codes are entered in binary, octal, or decimal.

Screen 4/5 (Puncturing and BER Options)

See [Figure 3-4](#). The parameter descriptions for this screen follow.

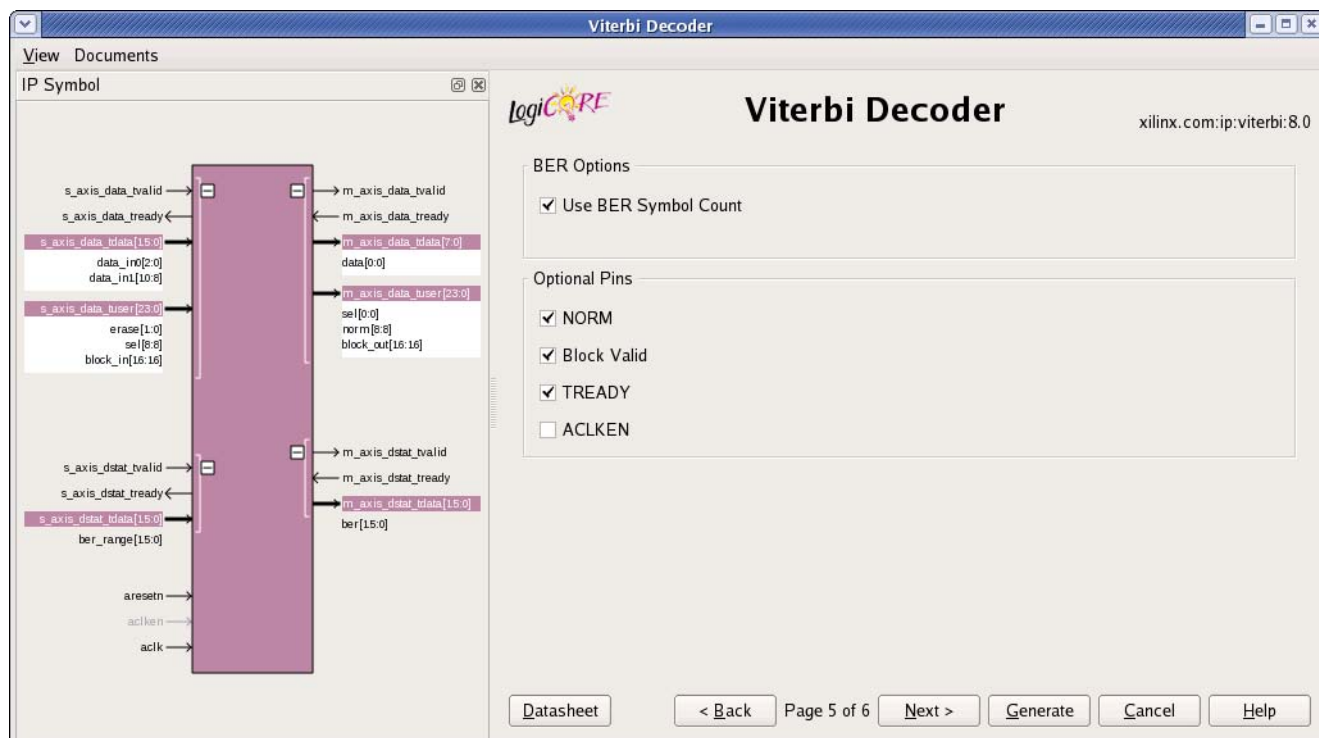


Figure 3-4: Viterbi Decoder Screen 5 (BER and Optional Pins)

BER Options

- Use BER Symbol Count:** Check this box if a Bit Error Rate (BER) monitor is required. The core compares the delayed incoming data with an encoded version of the outgoing decoded data to obtain an estimate of the BER on the channel. See [BER](#), page 32 for additional details. The Number of BER symbols over which the estimate is given is dynamically variable and is input on the `s_axis_dstat_tdata` bus and can range from 3 to $(2^{16}-1)$.

Optional Pins

Check the boxes of the optional pins that are required: NORM, BLOCK VALID, TREADY and ACLKEN. Select only pins that are genuinely required, because each selected pin results in more FPGA resources being used and can result in a reduced maximum operating frequency.

NORM

Check this box if a normalization output is required. For additional details, see [Normalization](#), page 32.

Block Valid

Check this box if BLOCK_IN and BLOCK_OUT signals are required. These signals track the movement of a block of data through the decoder. BLOCK_OUT corresponds to BLOCK_IN delayed by the decoder latency.

Parameter Ranges

Valid ranges for the parameters are shown in [Table 3-1](#).

Table 3-1: Parameter Ranges

Parameter	Min	Max	Notes
Channels	1	32	
Traceback Length	12	128	[1]
Constraint Length	7	9	-
Best State Width	3	8	-
Soft Width	3	5	-
Output Rates	2	7	[2]
Convolution Code	Bit width = constraint length		-

Notes:

- Traceback length must be divisible by 6 for the reduced latency case and ranges from 12 to 126.
- For the trellis mode, the output rate is always 2.

Parameter Values in the XCO File

Table 3-2: Parameter Values in the XCO File

GUI Name	Default Value	Valid Range	XCO parameter
Component Name	Viterbi_v8_0		component_name
ACLKEN	false	False/true	aclken
Architecture	Parallel	Parallel/Serial	architecture
BER Symbol Count	true	False/true	ber_symbol_count
Best State	true	False/true	best_state
Best State Width	3	3 to 8	best_state_width
Block Valid	false	False/true	block_valid
Channels	1	1 to 32	channels
Coding	Soft_Coding	Soft_Coding Hard_Coding	coding
Constraint Length	7	7 to 9	constraint_length
Convolution0 Code0	1111001	Number of bits must equal constraint length	convolution0_code0
Convolution0 Code1	1011011		convolution0_code1
Convolution0 Code2	0		convolution0_code2
Convolution0 Code3	0		convolution0_code3
Convolution0 Code4	0		convolution0_code4
Convolution0 Code5	0		convolution0_code5

Table 3-2: Parameter Values in the XCO File (Cont'd)

GUI Name	Default Value	Valid Range	XCO parameter
Convolution0 Code6	0		convolution0_code6
Convolution1 Code0	1111001		convolution1_code0
Convolution1 Code1	1011011		convolution1_code1
Convolution1 Code2	0		convolution1_code2
Convolution1 Code3	0		convolution1_code3
Convolution1 Code4	0		convolution1_code4
Convolution1 Code5	0		convolution1_code5
Convolution1 Code6	0		convolution1_code6
Convolution Code0 Radix	Binary	Binary Octal Decimal	convolution_code_0_radix
Convolution Code1 Radix	Binary	Binary Octal Decimal	convolution_code_1_radix
Data Format	Signed_Magnitude	Signed_Magnitude Offset_Binary	data_format
NORM	False	False/true	Norm
Output Rate0	2	2 to 7	output_rate0
Output Rate1	2	2 to 7	output_rate1
Puncturing	None	None External	puncturing
Reduced Latency	False	False/true	reduced_latency
Soft Width	3	3 to 5	soft_width
Traceback Length	42	12 to 128	traceback_length
TREADY	True	False/true	tready
Viterbi Type	Standard	Standard Multi_Channel Trellis_Mode Dual_Decoder	viterbi_type

Output Generation

Several files are produced when a core is generated, and customized instantiation templates for Verilog and VHDL design flows are provided in the .veo and .vho files, respectively. For detailed instructions, see the CORE Generator software [documentation](#)

Designing with the Core

This chapter includes guidelines and additional information to make designing with the core easier.

Functional Description

This core implements a Viterbi Decoder for decoding convolutionally encoded data. For details of the encoding process, see [Ref 1]. This is available in Xilinx® CORE Generator™ software. The decoder core consists of two basic architectures: a fully parallel implementation which gives fast data throughput at the expense of silicon area and a serial implementation which occupies a small area but requires a fixed number of clock cycles per decoded result.

Viterbi decoding decodes the data originally input to the convolutional encoder by finding an optimal path through all the possible states of the encoder. For a constraint length 7, there are 64 states and for a constraint length 9 there are 256. The basic decoder core consists of three main blocks as shown in Figure 4-1.

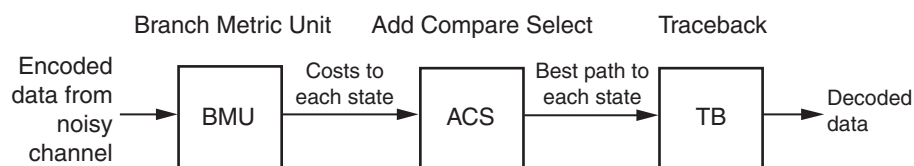


Figure 4-1: Viterbi Decoder Block Diagram

Costing

The first block is the branch-metric-unit (BMU). This module costs the incoming data. For the fully parallel decoder, the incoming data can be hard coded with bit width 1 or soft coded with a parameterizable bit width which can be set to any value from 3 to 5. Hard-coded data is decoded using the Hamming method of decoding, whereas soft data is decoded using an Euclidean metric. In hard coding, the demodulator makes a firm or hard decision on whether a one or zero is transmitted and provides no other information to the decoder on how reliable the decision is. For soft decoding, the demodulator provides the decoder with some side information together with the decision. The extra information provides the decoder with a measure of confidence for the decision. Soft-coded data gives a significantly better BER performance compared with hard-coded data. Soft decision offers approximately a 3 dB increase in coding gain over hard-decision decoding. Hard coding is available only for the Standard parallel or Multichannel Viterbi. Erasure (external puncturing) is not available with hard coding.

There are two available data formats for soft coding: soft signed magnitude and offset binary. See Table 4-1 for the data formats for the case of soft width 3.

Table 4-1: Data Format for Soft Width 3

	Signed Magnitude	Offset-Binary
Strongest 1	111	111
	110	110
	101	101
Weakest 1	100	100
Weakest 0	000	011
	001	010
	010	001
Strongest 0	011	000

Decoding

The second block in the decoder is the add-compare-select (ACS) unit. This block selects the optimal path to each state in the Viterbi trellis. Figure 4-2 shows one stage in the Viterbi trellis for a constraint length 3 decoder. The ACS block uses the convolutional codes to extract the correct cost from the BMU for each branch in the trellis.

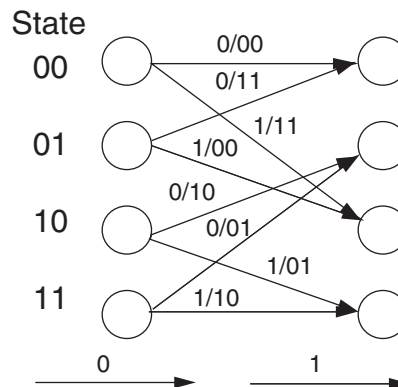


Figure 4-2: State Transitions for Constraint Length 3

The BER performance of the Viterbi algorithm varies greatly with different convolutional code sets; some of the standard convolutional codes are shown in Table 4-2.

Table 4-2: Standard Convolution Codes

Constraint Length	Output Rate = 2		Output Rate = 3	
	binary	octal	binary	octal
7	1111001 1011011	171 133	1001111 1010111 1101101	117 127 155
9	101110001 111101011	561 753	101101111 110110011 111001001	557 663 711

The ACS module decodes for each state in the trellis. Thus, for a constraint length 7 decoder which has 64 states, there are 64 sub-blocks in the ACS block.

If the core is implemented in serial mode, the amount of silicon required for each sub-block is very small. A decoder of constraint length 7 can be implemented on a small Spartan[®] device. See [Performance Characteristics, page 5](#) for further characterization of the decoder.

Traceback

The final block in the decoder is the traceback block. The actual decoding of symbols into the original data is accomplished by tracing the maximum likelihood path backwards through the trellis. Up to a limit, a longer sequence of tracing results in a more accurate path through the trellis. After a number of symbols equal to at least six times the constraint length, the decoded data is output. Thus the Viterbi always requires at least twice the traceback length of data to be input subsequently to a given input for that input to be successfully decoded. The traceback starts from zero or best state; the best state is estimated from the ACS costs. The traceback length is the number of trellis states processed before the decoder makes a decision on a bit. The decoded data is output only after a traceback length number of bits has been traced through. In other words, the traceback length determines the length of the training sequence for the Viterbi Decoder.

The length of the traceback is parameterizable and can be set to any value between 12 and 128. For the reduced latency option, the traceback length can only be a multiple of 6 between 12 and 126. The recommended value for non-punctured decoding is at least 6 times the constraint length. For data that has been punctured, that is, symbols removed prior to transmission over the channel, a larger value traceback length is required; usually, it is at least 12 times the constraint length to obtain optimal BER performance.

A best state option is available to select the starting location for the traceback from the state with minimal cost. There is also a reduced latency option which reduces the latency on the core by approximately half. The latency is of the order of two times the traceback length for reduced latency; see [Latency, page 6](#). The reduced latency option has a slight speed penalty and is not available with the multichannel or serial core. The traceback block is implemented in block RAM, and the larger the value of the traceback length, the greater the block RAM requirements. See [Performance Characteristics, page 5](#) for additional details.

Serial Decoder

Data can be processed a bit at a time if the serial option is selected. This results in a smaller design but also a significant increase in latency. The number of clock cycles needed to process each set of input symbols depends on the output rate and the soft width of the data. See [Table 4-3](#) for the minimum number of clock cycles required for each set of input symbols.

$$\text{number of clock cycles} = \text{soft_width} + \text{output_rate} + 6$$

The enabling of the data through the core is controlled by the `tvalid` and `tready` signals on the `s_axis_data_tdata`. See [Figure 4-3](#) for an example of the AXI handshaking on the serial core.

Table 4-3: Minimum Required Clock Cycles Per Input Symbol Set for a Rate 1/2 Serial Decoder

Soft Width	Minimum Clock Cycles
3	11
4	12
5	13

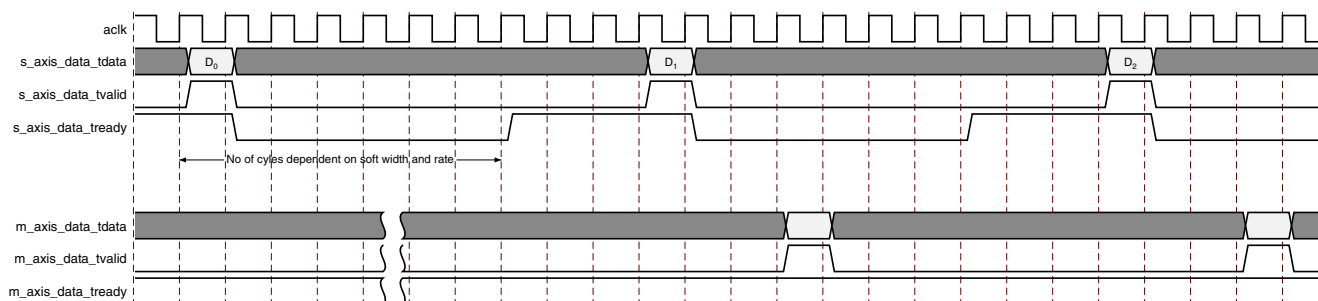


Figure 4-3: Handshaking Example on the Serial Core

Multichannel Decoder

The multichannel decoder decodes many interlaced channels using a single Viterbi Decoder. The input to the multichannel decoder is interlaced encoded data on the slave DATA channel. For a channel count of 3, channel 1 data is input followed by channel 2 and then channel 3 in a repeating sequence. The output is interlaced decoded data on the master DATA channel. See Figure 4-4. The multichannel decoder can decode from 2 to 32 channels. The larger the number of channels, the greater the block RAM requirements, as each channel requires its own traceback.

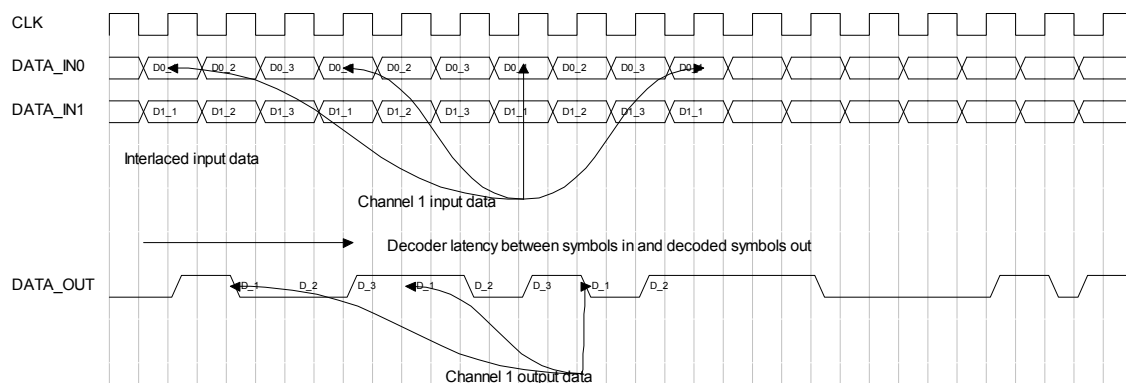


Figure 4-4: Multichannel Decoder with Channel Count 3

The BER from the multichannel decoder, if selected, gives the average number of errors present over all the input channels.

The multichannel decoder can decode at high speeds, but the true output rate of the decoder is equal to the speed divided by the number of channels. See [Performance Characteristics](#), page 5 for characterization results on the multichannel decoder.

Trellis Mode Decoder

For systems that are both power-limited and bandwidth-limited, Trellis Coded Modulation (TCM), or Pragmatic Trellis Coded Modulation (PTCM) as it can be known, is used. The modulation schemes are generally 8-PSK and 16-PSK. The LogiCORE™ IP Viterbi Decoder as used in a trellis decoder system is shown by the grayed out box in the trellis mode system diagram shown in Figure 4-5. In trellis mode, the BMU in the Viterbi is not used, but the costing is done externally using a Branch Metrics cost table. The ACS and traceback sections are used as normal. The output from the cost table is four TCM buses and a 4-bit Sector Bus. The address for the Branch Metrics cost table is provided by the I and Q outputs after symbol recovery from the PSK demodulator.

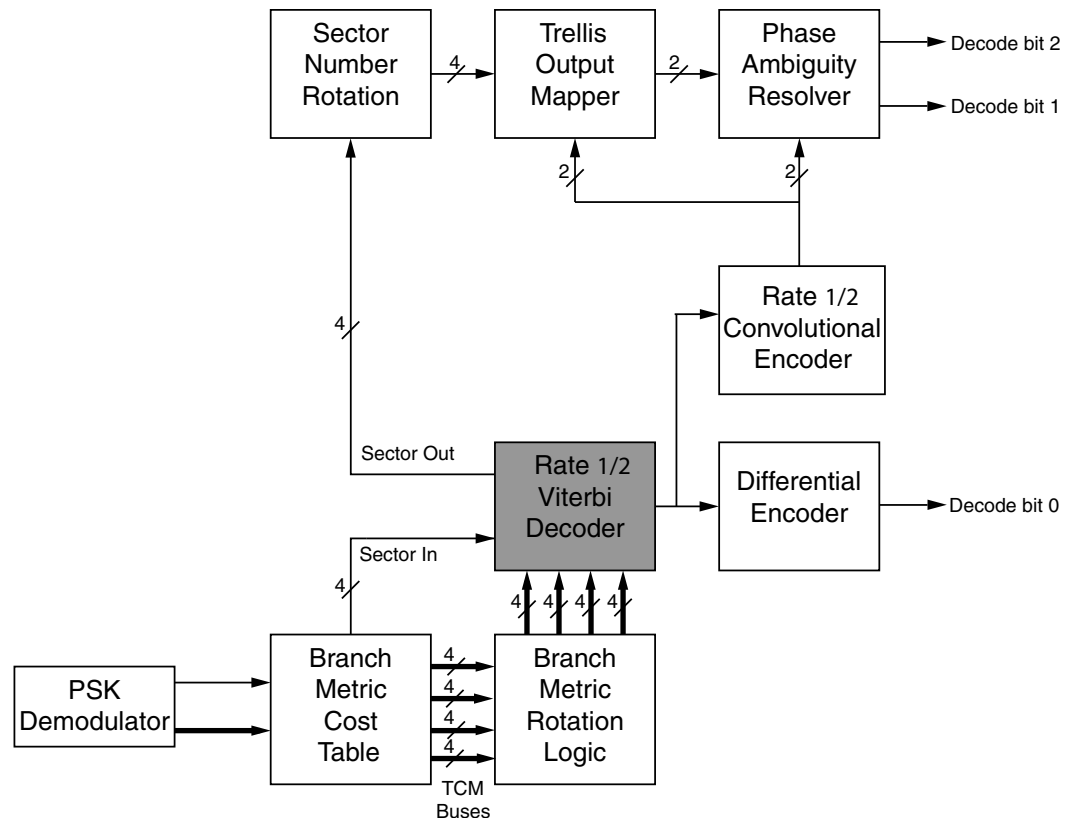


Figure 4-5: Viterbi Decoder in Trellis Receiver

Pragmatic Trellis Coded Modulation (PTCM) has the following setup:

- A standard rate 1/2 Viterbi Decoder is used.
- The data is costed externally to the decoder, bypasses the BMU, and sent to the ACS.

Some points to note:

- Each TCM bus can be 4 to 6 bits wide depending upon the TCM bus width selected on the Viterbi Decoder.
- The 4 to 6 bits are unsigned cost values that are applied to the ACS module in the Viterbi Decoder. If the width of the generated costs is less than the (soft width+1), then the data should be tied to the lower bits and the remaining TCM input bits tied to zero.
- The received symbol or phase angle is converted to four branch metrics and a sector number externally to the decoder using an external lookup table.

- The sector number identifies the part of the I-Q plane where the symbol was received. The branch metrics are then processed by the Trellis Mode Decoder. The sector number is delayed by the Viterbi latency in the Trellis Mode Decoder.
- The costed data and sector are input to the core on the `s_axis_data_tdata` bus. The decoded data along with the output sector are output from the core on the `m_axis_data_tdata` bus.

Dual Rate Decoder

For a given constraint length and traceback-length, the core can function as a Dual Decoder, that is, two sets of convolutional codes and output rates can be used internally to the decoder. The dual-decoder offers significant device area savings when two different decoders with the same constraint length are required. For example, as a constraint length 7 decoder the core can decode as a rate 1/2 decoder and a rate 1/3 decoder. The implementation requires only a little additional logic for the extra costing involved in the BMU and some multiplexing in the ACS unit (see Figure 4-1). The Dual Decoder can be implemented as either parallel or serial architecture, and erasure pins can be present on the input. The selection of the decoder rate and codes is through the SEL pin (see Figure 4-6). When the SEL pin is low, output rate0 and convolution0_codes are used in the decoding. When the SEL pin is high, then the rate is 1/output_rate1, and the convolution1_codes are used to decode the incoming data. The SEL_O pin shows the decoded data corresponding to the original SEL set of data points. The number of input data buses is equal to the max of the two output rates.

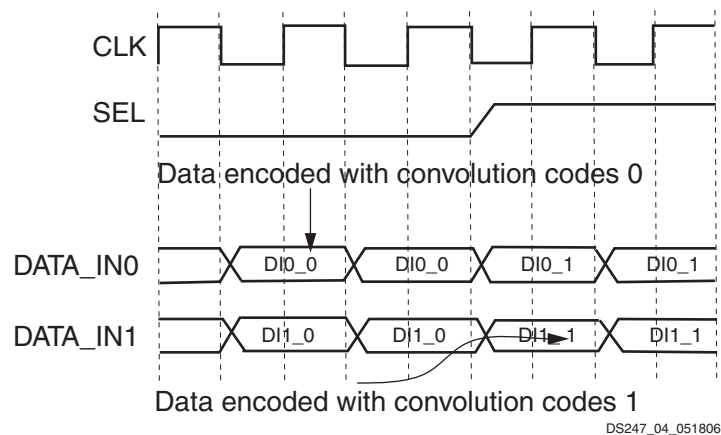


Figure 4-6: SEL Input for Dual Decoder

Erasure

If the data has been punctured prior to transmission, then de-puncturing is carried out externally to the Viterbi Decoder (see Figures 4-7 and 4-8). The presence of null-symbols (that is, symbols which have been deleted prior to transmission across the channel) is indicated using the erasure input ERASE. The decoder functions exactly as a non-punctured Viterbi Decoder, except the corresponding DATA_IN inputs are ignored when the erased input bit is high. If ERASE(0) is high, then the input on DATA_IN0 is viewed as a null-symbol. If ERASE(1) is high, then the input on DATA_IN1 is viewed as a null symbol, etc. Although the normal usage of erasure is with a rate 1/2 decoder, the erasure pins can be present for output rates greater than 2. Erasure can be used with the

Standard, Multichannel and Dual Decoder. Erasure cannot be present on the Trellis Mode Decoder. See the timing diagram in Figure 4-9 for a decoder working with external rate 3/4 erasure. Note that DATA_IN0, DATA_IN1 are the fields on the s_axis_data_tdata bus and the ERASE(0) and ERASE(1) are on the s_axis_data_tuser bus. The output data, DATA_OUT is on the m_axis_data_tdata bus.

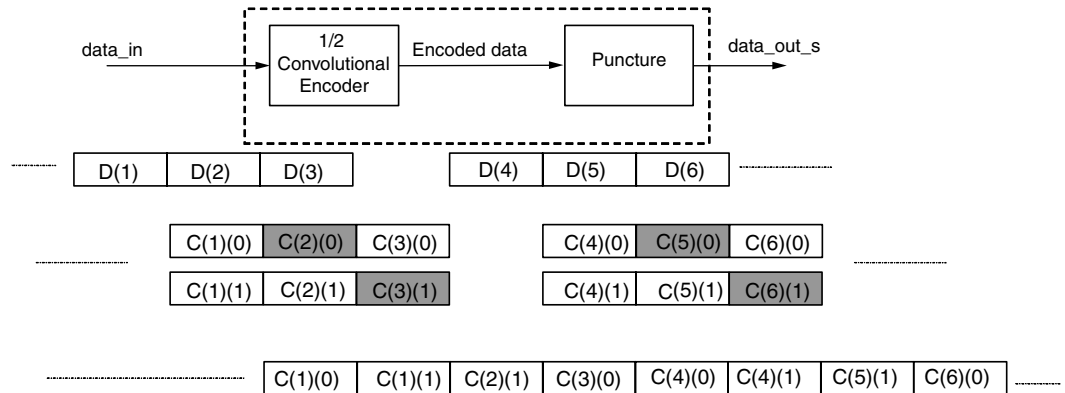


Figure 4-7: Puncturing Encoded Data with 3/4 Puncture Rate with Single-Channel Output

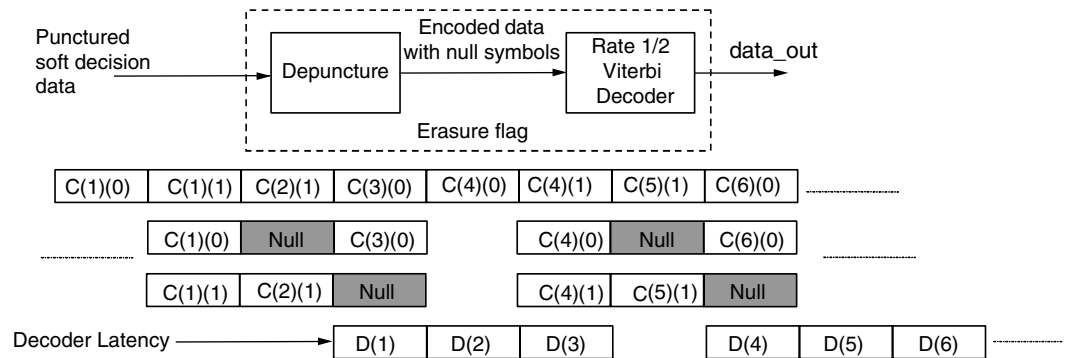


Figure 4-8: De-puncturing Rate 3/4 Punctured Data with Single-Channel Soft Input and Erasure

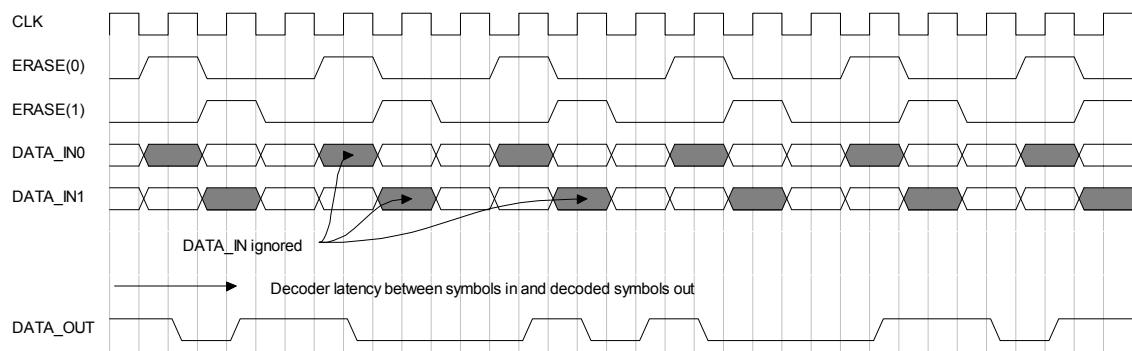


Figure 4-9: Viterbi Decoder with Erasure Input Following a Rate 3/4 Pattern

BER

The Bit Error Rate (BER) option on the decoder monitors the error rate on the transmission channel. Decoded data from the Viterbi Decoder is re-encoded using the convolutional encoder and compared with a delayed version of the data input to the decoder. An error is indicated if the delayed and encoded data differ (see Figure 4-10). The count is incremented on a symbol-by-symbol basis. For example, if both the I & Q outputs differ from the expected I and Q for a rate 1/2 decoder, then this is only considered as one error on the BER count. The two sets of symbols can differ if there is an error on the channel or if the Viterbi Decoder has decoded incorrectly. The probability of the decoder incorrectly decoding is significantly smaller than the probability of a channel bit error; therefore the BER output gives a good estimate of the errors on the channel.

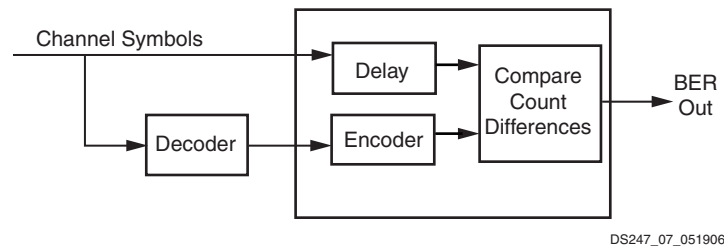


Figure 4-10: Bit Error Rate Calculation

The BER symbol count which is input on the DSTAT channel determines the number of input symbols over which the error count takes place. This value can be varied dynamically. The output error count BER is the number of errors that has been counted during the BER symbol count. `m_axis_dstat_valid` is asserted when BER symbol count input symbols have been processed and the bit error rate value is present on `m_axis_dstat_tdata`. The BER output always has a width of 16; therefore the maximum number of errors that can be counted is $(2^{16}-1)$. If more errors occur than this upper limit, the maximum number of errors is output, that is, BER is set to all 1s.

Normalization

The NORM signal is an optional output that gives immediate monitoring of the errors on the channel. As the metrics grow in the ACS unit, they must be normalized to avoid overflow. When normalization occurs the decoder subtracts a fixed value from all metrics and asserts the normalization signal. If the ACS unit requires normalization, then there are uncertainties or errors on the channel. The more frequent the normalization, the higher the rate of errors present. The actual frequency of the normalization depends on many factors, in particular the soft width and the output rate of the decoder.

The normalization signal can be used to detect Viterbi synchronization. A high normalization rate (exceeding a certain threshold) indicates loss of synchronization. A low normalization rate (less than a certain threshold) indicates Viterbi synchronization has been achieved. In general, the normalization rate is inversely proportional to the Signal-to-Noise Ratio (SNR) at the decoder input.

Synchronization

If required the synchronization status of the core can be monitored externally to the core. The method involves the analysis of both the normalization output from the ACS modules within the core and the BER performance of the core. A complete description of the method used is provided in [Ref 8].

The Normalization rate by itself gives an indication of the Viterbi Decoder synchronization status. A high normalization rate, exceeding a predetermined threshold, implies a loss of synchronization. Similarly, the BER rate of the core, by itself, can indicate a loss of synchronization. Thus, if the BER rate, or the normalization rate, was used in isolation, the threshold required would vary with the noise on the channel. The BER rate, like the normalization rate, would therefore be dependent on the signal to noise ratio E_b/N_o . To monitor synchronization, both the BER rate and the normalization rate are required to achieve a synchronization method that is independent of E_b/N_o .

Packet Handling

Viterbi decoding is a continuous operation, but the input to the decoder can be packet based rather than continuous streams. For data encoded in packets, it is necessary to terminate the encoder between the packets by the insertion of what is called zero tail bits. For a constraint length 7 decoder, there are 6 zero bits inserted into the encoder at the end of the packet. For a general Viterbi (constraint length -1), tail bits are required to return the encoder back to state zero. The effect of these zero tail bits is to return the Viterbi trellis to zero state and also the next packet starts from state zero. The Viterbi handles tail bits when working in any of the basic modes; no additional signals or control circuitry is required.

Although zero-tail bits or zero-tail termination is the standard method for handling packets within the Viterbi Decoder, there is a rate loss on the channel caused by constraint length -1 information bits being added to the original message. If the original packet contained m bits, the output code words are of length $m + K - 1$, where K is the constraint length. Thus, the effective rate on the channel becomes:

$$\text{RateNew} = \text{Rate} \times \left(1 - \frac{K - 1}{m + K - 1}\right)$$

For large packets, the rate loss becomes insignificant. For smaller packets, the method of tail-biting avoids the issue of the fractional rate loss by letting the last $K-1$ information bits define the starting state of the encoder. Only the data is encoded, that is, exactly m encoded bits are produced. In this case, no rate loss occurs and the encoding always starts and ends in the same state, but not necessarily the zero state. For a full description of the Viterbi Decoder and trellis termination and tail-biting, see [Ref 7].

Design Guidelines

The following section provide guidelines on the best design practices for inclusion of the LogiCORE™ IP Viterbi Decoder core within a system.

Data Format

Use soft coding, as this gives a better BER performance. For the soft coding width, 3 or 4 bits should be sufficient. If the width is any larger, the size of the ACS is increased but does not give a large improvement in BER performance.

The format for soft width 3 is shown in Table 4-1. To extend the soft width input, for example, 4 bits, see the example shown:

	Signed Magnitude		Offset Binary	
Strongest..1	1111	1111
Weakest..1	1000	1000
Weakest..0	0000	0111
Strongest..0	0111	0000

For the signed magnitude, it is the lower bits that are extended. The MSB is the sign bit. For offset binary, a lower bit is added to extend the range. In the example, this is a range from low, 0, to high, 15.

Data Input

Remove the DC bias in the soft input data before entering the data into the Viterbi Decoder. Use symmetric rounding when the soft data width is larger than the Viterbi input data width. Note that the Viterbi input format is a balanced number system, while the 2's complement number system is unbalanced with one more negative number than positive number, which can lead to 1/2 LSB DC offset if mapping does not take this into consideration.

Best State

This is on by default because it tends to give a better BER performance, especially for punctured data. The improvement is in the order of 0.25 dB. For non-punctured data, the option is not really needed. However, there is a penalty in the area of the core if this option is used. There is the flexibility to change the width if there are issues with noise.

Control Signals

Use the `aresetn` so as to put the decoder in a good start-up state. Make use of the DCM Lock signal to control the `aresetn`. If there is a need to save power, and not use the core, use the `aclken` to stop, or start the core operations. However, do not use `aclken` or `aresetn` to deal with packets of data. If dealing with blocks of data, use the BlockIn/Out signals. To qualify the output data AND the `tvalid` and BlockOut signals together.

TVALID

Use the `tvalid` signal to validate the data. This signal could be used as a clock enable to store the output from the Viterbi Decoder, for example, in a FIFO or memory.

NORM

Use the NORM signal to get an idea of the errors in the Viterbi Decoder and how it is dealing with the internal path metrics. If it is a noisy system, there will be normalizing of the metrics and a NORM pulse will appear.

Keep it Registered

To simplify timing and increase system performance in an FPGA design, keep everything registered, that is, all inputs and outputs from the user application should come from, or connect to a flip-flop. While registering signals might not be possible for all paths, it simplifies timing analysis.

Viterbi Decoder Non-features Summary

This section outlines the features that might be external to the core, depending upon the system being implemented.

Multichannel BER

If there is a requirement for individual BER circuits, these have to be done externally to the core.

Log Likelihood Ratio

The Log Likelihood Ratio (LLR) values that create the soft code inputs have to be generated externally to the core. See resources on the internet for information on how to do this.

De-puncturing

The Viterbi Decoder does not perform any kind of internal de-puncturing. Any de-puncturing has to be performed externally.

Trellis Mode

In this mode, the external costing table has to be created. See [Trellis Mode Decoder](#).

Factors Affecting BER Performance

When viewing the Viterbi Decoder as part of a whole communication system, there are a few parameters that could affect BER performance. For the core, parameters such as data format, constraint length, and traceback length, use of best state have an effect. However, these are generally determined by the standard being implemented. After these have been taken care of, then outside factors need to be taken into consideration, for example modulation type, the channel model used, E_b/N_0 value, convolutional codes and DC bias. For factors that affect packets of data, see [\[Ref 7\]](#).

Detailed Example Design

Demonstration Test Bench

When the core is generated using CORE Generator™, a demonstration test bench is created. This is a simple VHDL test bench that exercises the core.

The demonstration test bench source code is one VHDL file: `<component_name>/demo_tb/tb_<component_name>.vhd` in the CORE Generator output directory. The source code is comprehensively commented.

Using the Demonstration Test Bench

The demonstration test bench instantiates the generated Viterbi Decoder core. If the CORE Generator project options were set to generate a structural model, a VHDL or Verilog netlist named `<component_name>.vhd` or `<component_name>.v` was generated. If this file is not present, generate it using the netgen program, for example:

```
netgen -sim -ofmt vhdl <component_name>.ngc <component_name>.vhd
```

Compile the netlist and the demonstration test bench into the work library (see your simulator documentation for more information on how to do this). Then simulate the demonstration test bench. View the test bench's signals in your simulator's waveform viewer to see the operations of the test bench.

The Demonstration Test Bench in Detail

The demonstration test bench performs the following tasks:

- Instantiates the core
- Generates a clock signal
- Generates a source data table consisting of a sinusoid
- Serializes and convolution encodes the source data to create input data for the Viterbi Decoder core
- Inserts errors into the input data to demonstrate error correction
- If external erasure is supported, erase inputs using a predefined puncture pattern
- Drives the core's input signals to demonstrate core features
- Checks that the core's output signals obey AXI4 protocol rules (data values are not checked in order to keep the test bench simple)
- Provides signals showing the separate fields of AXI4-Stream TDATA and TUSER signals
- Provides signals showing the source data before serialization and encoding, and the deserialized decoded output data

The demonstration test bench drives the core input signals to demonstrate the features and modes of operation of the core. The operations performed by the demonstration test bench are appropriate for the configuration of the generated core and are a subset of the following operations:

1. An initial phase where the core is initialized and no operations are performed.
2. Decode data containing no errors.
3. Decode and correct data containing some errors.
4. Try to decode and correct data containing many errors, sometimes failing to correct the errors.
5. If BER statistics are supported, set up BER statistics over various ranges.
6. Demonstrate the use of AXI4-Stream handshaking signals `tvalid` and `tready`.
7. Demonstrate the effect of asserting `aresetn`.
8. If `ac1ken` is present: demonstrate the effect of toggling `ac1ken`.

Customizing the Demonstration Test Bench

It is possible to modify the demonstration test bench to use different source data or different control information.

Source data is pre-generated in the `create_src_table` function and stored in the `SRC_DATA` constant. Data from this constant is serialized by the `s_data_stimuli` process, convolution encoded by the `conv_data` procedure and driven into the core by the `encode_data` procedure, which also inserts errors and controls erasures. Data is driven continuously throughout the operation of the test bench: new input data is required for the Viterbi Decoder core to produce output data.

Source data before serialization is shown on the `s_axis_data_tdata_src_des` signal, which is synchronized to the corresponding serialized and encoded data being driven into the Viterbi Decoder core. Output data is deserialized by the `deserialize_output` process and shown on the `m_axis_data_tdata_des` signal. These signals can be viewed in a simulator to compare source data and decoded, corrected output data.

If external erasure is supported, the test bench models the effect of puncturing by a convolution encoder, by de-puncturing and inserting null symbols in the input data. The puncture rate is defined by the `PUNC_INPUT_RATE` and `PUNC_OUTPUT_RATE` constants, and the puncture pattern is defined in the `PUNC_CODES` constant. These constants can be modified to change the puncture rate and pattern.

BER statistics are controlled by the `s_dstat_stimuli` process: this selects the BER range and drives transactions on the `S_AXIS_DSTAT` channel to set up BER statistics generation.

The clock frequency of the core can be modified by changing the `CLOCK_PERIOD` constant.

Migrating

This appendix describes migrating from older versions of the IP to the current IP release.

Parameter Changes in the XCO File

The CORE Generator™ core update functionality can be used to update an existing XCO file from v7.0 to v8.0, but the update mechanism alone does not create a core compatible with v7.0. [Table A-1](#) shows the changes to XCO parameters from v7.0 to v8.0.

Table A-1: Parameter Changes in XCO File

Version 7.0	Version 8.0	Notes
architecture	Unchanged	
ber_symbol_count	Ber_symbol_count	Different functionality as the BER count is now input on dynamically on the DSTAT channel
ber_thresh	Removed	Synchronization not supported in v8.0
best_state	Unchanged	
best_state_width	Unchanged	
block_valid	Unchanged	
ce	acklen	Renamed from ce to acklen for AXI4 standardization
channels	Unchanged	
coding	Unchanged	
component_name	Unchanged	
constraint_length	Unchanged	Support constraint length reduced to 7 to 9
convolution0_code0	Unchanged	
convolution0_code1	Unchanged	
convolution0_code2	Unchanged	
convolution0_code3	Unchanged	
convolution0_code4	Unchanged	
convolution0_code5	Unchanged	
convolution0_code6	Unchanged	

Table A-1: Parameter Changes in XCO File (Cont'd)

Version 7.0	Version 8.0	Notes
convolution1_code0	Unchanged	
convolution1_code1	Unchanged	
convolution1_code2	Unchanged	
convolution1_code3	Unchanged	
convolution1_code4	Unchanged	
convolution1_code5	Unchanged	
convolution1_code6	Unchanged	
convolution_code_0_radix	Unchanged	
convolution_code_1_radix	Unchanged	
data_format	Unchanged	
direct_traceback	Removed	Packet options not supported in v8.0
dynamic_thresholds	Removed	Synchronization not supported in v8.0
maximum_direct	Removed	Packet options not supported in v8.0
norm	Unchanged	
norm_thresh	Removed	Synchronization not supported in v8.0
number_of_ber_symbols	Removed	BER symbol count is input dynamically
output_rate0	Unchanged	
output_rate1	Unchanged	
puncturing	Unchanged	
rdy		Replaced with AXI4 control signals
reduced_latency	Unchanged	
soft_width	Unchanged	
synchronization	Removed	Synchronization not supported in v8.0
synchronous_clear	aresetn	aresetn always present on core
traceback_length	Unchanged	
trellis_initialization	Removed	Packet options not supported in v8.0
viterbi_type	Unchanged	

Table A-2: Port Changes

Version 7.0	Version 8.0	Notes
DATA_IN0	s_axis_data_tdata	Now exists as a field within s_axis_data_tdata
DATA_IN1	s_axis_data_tdata	
DATA_IN2	s_axis_data_tdata	
DATA_IN3	s_axis_data_tdata	
DATA_IN4	s_axis_data_tdata	
DATA_IN5	s_axis_data_tdata	
DATA_IN6	s_axis_data_tdata	
TCM00	s_axis_data_tdata	Now exists as a field within s_axis_data_tdata when the core is in trellis mode
TCM01	s_axis_data_tdata	
TCM10	s_axis_data_tdata	
TCM11	s_axis_data_tdata	
SECTOR_IN	s_axis_data_tdata	Occupies the top byte of the input tdata when the core is in trellis mode
BLOCK_IN	s_axis_data_tuser	Now exists as a field within s_axis_data_tuser
PACKET_START	Removed	Packet processing removed in v8.0 of the core
TB_BLOCK	Removed	
PS_STATE	Removed	
TB_STATE	Removed	
ber_thresh	Removed	Synchronization removed in v8.0 of the core
NORM_THRESH	Removed	
ERASE	s_axis_data_tuser	Now exists as a field within s_axis_data_tuser
DATA_OUT	m_axis_data_tdata	Now exists as a field within m_axis_data_tdata
DATA_OUT_DIRECT	Removed	Packet processing removed in v8.0 of the core
DATA_OUT_REVERSE	Removed	
PACKET_START_O	Removed	
TB_BLOCK_O	Removed	
DIRECT_RDY	Removed	
REVERSE_RDY	Removed	
BER	m_axis_dstat_tdata	The dstat input and output channel now handle the BER count which can be varied dynamically
BER_DONE	m_axis_dstat_tvalid	

Table A-2: Port Changes (Cont'd)

Version 7.0	Version 8.0	Notes
NORM	m_axis_data_tuser	Now exists as a field within m_axis_data_tuser
SECTOR_OUT	m_axis_data_tdata	Occupies the top byte of the output data if the core is in trellis mode
BLOCK_OUT	m_axis_data_tuser	Now exists as a field within m_axis_data_tuser
OUT_OF_SYNC	Removed	Synchronization removed in v8.0 of the core
OOS_FLAG	Removed	
SEL	s_axis_data_tuser	Now exists as a field within s_axis_data_tuser
SEL_O	m_axis_data_tuser	Now exists as a field within m_axis_data_tuser
ND	s_axis_data_tvalid	
RFD	s_axis_data_tready	
RDY	m_axis_data_tready	
CE	aclken	Rename only
SCLR	aresetn	Rename and change on sense (now active Low). Must now be asserted for at least 2 cycles
CLK	aclk	Rename only

Debugging

If a Viterbi Decoder is not functioning as expected, here are some tips to consider. Xilinx [Technical Support](#) can also be contacted.

1. See the examples in the demonstration test bench to see if they match your configuration.
2. If not, create a simple design based on your parameters and one of the demonstration test bench examples.
3. Check that parameters of the encoder and the Viterbi Decoder agree for code rate and convolutional codes. Confirm that the codes are tied up correctly, that is, that you are not applying convolution code0 to convolution code1 or similar.
4. Check that the soft width data format is set up correctly, that is, signed or offset, and that the data format for the data input to the decoder is correct, for example, the MSB on the data input is the sign bit for signed magnitude.
5. Add BER and NORM ports to monitor errors.
6. Run the decoder in both functional simulation and post-PAR simulation.
7. To speed testing, consider the use of a ChipScope™ analyzer or HW Cosim to find the errors.

If a ChipScope analyzer is used, monitor the following signals:

- Data Inputs
- ARESETN
- TVALID
- TUSER (NORM)
- DSTAT for BER count

For punctured code, ensure:

- Erase input is used

For packet data:

- TUSER (Block In)
- TUSER (Block Out)

8. Check the simulation.

If only functional simulation is working correctly, then check timing simulation. See [\[Ref 6\]](#).

Additional Resources

Xilinx Resources

For support resources such as Answers, Documentation, Downloads, and Forums, see the Xilinx Support website at:

<http://www.xilinx.com/support>.

For a glossary of technical terms used in Xilinx documentation, see:

http://www.xilinx.com/support/documentation/sw_manuals/glossary.pdf.

Solution Centers

See the [Xilinx Solution Centers](#) for support on devices, software tools, and intellectual property at all stages of the design cycle. Topics include design assistance, advisories, and troubleshooting tips.

References

1. Convolutional Encoder Product Guide, [PG026](#)
2. Convolutional Encoder [Product Page](#)
3. Viterbi Decoder [Product Page](#)
4. Xilinx AXI Design Reference Guide [UG761](#)
5. [AMBA 4 AXI4-Stream Protocol Version: 1.0 Specification](#)
6. Synthesis and Simulation Design Guide ([UG626](#))
7. Viterbi Decoder Block Decoding - Trellis Termination and Tail-Biting ([XAPP551](#))
8. Viterbi Synchronization ([DS205](#))

Technical Support

Xilinx provides technical support at www.xilinx.com/support for this LogiCORE™ IP product when used as described in the product documentation. Xilinx cannot guarantee timing, functionality, or support of product if implemented in devices that are not defined in the documentation, if customized beyond that allowed in the product documentation, or if changes are made to any section of the design labeled DO NOT MODIFY.

See the IP Release Notes Guide ([XTP025](#)) for more information on this core. For each core, there is a master Answer Record that contains the Release Notes and Known Issues list for the core being used. The following information is listed for each version of the core:

- New Features
- Resolved Issues
- Known Issues

Ordering Information

Contact your local Xilinx [sales representative](#) for pricing and availability of additional Xilinx LogiCORE IP modules and software. Information about additional Xilinx LogiCORE IP modules is available on the Xilinx [IP Center](#).

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
01/18/12	1.0	Initial Xilinx release as a Product Guide. Previous non-AXI Data Sheet/User Guide is DS247/UG745.

Notice of Disclaimer

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of the Limited Warranties which can be viewed at <http://www.xilinx.com/warranty.htm>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in Critical Applications: <http://www.xilinx.com/warranty.htm#critapps>.

© Copyright 2012 Xilinx, Inc. Xilinx, the Xilinx logo, Artix, ISE, Kintex, Spartan, Virtex, Zynq, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. All other trademarks are the property of their respective owners.