LZR1 Core

High-Performance Lossless Compression Core

General Description

Lossless data compression is a class of data compression algorithms that allows the exact original data to be reconstructed from the compressed data. Lossless compression is used when it is important that the original and the decompressed data be identical, or when no assumption can be made on whether certain deviation is uncritical. Typical applications include data storage and transmission.

LZR1 implements the lossless compression algorithm on short units of data ("frames"). The core supports frame sizes up to 4096 bytes.

The design is fully synchronous and available in multiple configurations varying in bus widths and throughput.

LZR1-6 can easily deliver 10 Gbps of throughput in both FPGA and ASIC implementations. The compression ratio greatly depends on the data and somewhat depends on the frames size; on typical file corpuses varies between 1.5 and 2.

Symbol

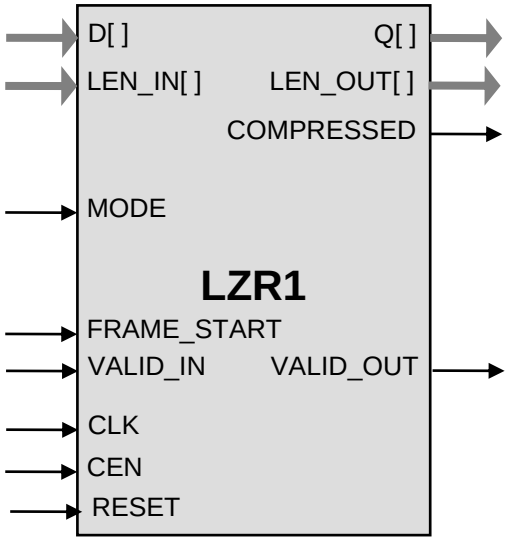

Key Features

Each frame is compressed and decompressed independently

High throughput: LZR1 easily scales to 10 Gbps in FPGA

Compatibility with public-domain LZ software implementations allows for interoperability

Support for compression and decompression in a single core; dedicates compression and decompression versions are available

Back-to-back compression with no gaps between the frames

Applications

- Disk and tape storage systems
- High-performance solid-state storage
- Networking, including cellular backhaul

Pin Description

Name	Type	Description
CLK	Input	Core clock signal
RESET	Input	HIGH level asynchronously resets the core
CEN	Input	Clock enable, bringing this input low pauses the core
MODE	Input	LOW level for compression, HIGH for decompression operation
FRAME_START	Input	Indicates the start of the new frame
D[ ]	Input	Input data bus
LEN_IN[ ]	Input	Input data length in bytes
VALID_IN	Input	HIGH if the FRAME_START begins the new frame to compress / decompress
Q[ ]	Input	Output data bus
LEN_OUT[ ]	Input	Output data length in bytes
VALID_OUT	Output	HIGH indicates a valid frame on the output
COMPRESSED	Output	HIGH indicates that the compression was successful

Function Description

The core implements lossless compression and decompression of short (up to 4096 bytes) frames of data. If the compression fails to produce a shorter result, the core copies the input frame to output and indicates the uncompressed status via a dedicated pin.

Performance

On the Calgary corpus with the 4096-byte block the core exhibits the following performance:

File	Compression ratio	File	Compression ratio
bib	1.56 (64.22%)	paper3	1.49 (66.93%)
book1	1.39 (71.97%)	paper4	1.54 (64.83%)
book2	1.58 (63.10%)	paper5	1.59 (62.97%)
geo	1.13 (88.52%)	paper6	1.64 (60.88%)
news	1.51 (66.05%)	pic	4.08 (24.52%)
obj1	1.60 (62.70%)	progc	1.74 (57.59%)
obj2	1.83 (54.53%)	progl	2.13 (46.84%)
paper1	1.61 (62.02%)	progp	2.13 (47.00%)
paper2	1.54 (65.11%)	trans	1.93 (51.73%)

Synthesis results

Configuration with 4096 bytes block in Virtex 5

	10 Gbps @ 156 MHz	
Compression/Decompression switchable	Slices	36 Kbit memories
	2200	36

Configuration with 4096 bytes block in TSMC 65 nm G+ process

	6 Gbps @ 250 MHz			6.4 Gbps @ 200 MHz		
	Gates	Memory		Gates	Memory	
		Dual Port 1K x 32	Single Port 4K x 12		Dual Port 1K x 32	Single Port 4K x 12
Compression only	35K	10	5	42K	12	6
Decompression only	30K	10	-	36K	12	-
Compression/Decompression switchable	45K	10	5	54K	12	6

Deliverables

HDL Source Licenses

- Synthesizable Verilog RTL source code
- Verilog testbench (self-checking)
- Vectors for the testbench
- Expected results
- User Documentation

Contact Information

IP Cores, Inc.
3731 Middlefield Rd.
Palo Alto, CA 94303, USA
Phone: +1 (650) 815-7996
E-mail: info@ipcores.com
www.ipcores.com