

Iterative Detection Read Channel Technology in Hard Disk Drives

Table of Contents

Executive Summary	1
Background.....	2
Achieving Enhanced Performance Through IDRC	3
Soft Information	3
Coding	3
Iterative Detector.....	5
Data Integrity	6
Looking Forward	6

Iterative Detection Read Channel Technology in Hard Disk Drives

Authors: Richard Galbraith and Travis Oenning

Executive Summary

Signal processing technology has made dramatic advances since the early 1990s. The development of iterative decoding schemes has influenced almost every form of communication technology, including satellite transmissions, cellular telephony and wireless networks. By leveraging the synergy between iterative decoding and the introduction of ultra-dense 65 nanometer (nm) System-On-Chip (SOC) technology to its next-generation hard drives, HGST, a Western Digital Company, is realizing key reliability enhancements and taking signal processing to a new level.

When adapted to hard drive design, the powerful iterative decoding approach developed by HGST is referred to as Iterative Detection Read Channel (IDRC) technology. This breakthrough approach is delivering enhancements to signal-to-noise ratios that are unprecedented over the last 20 years of hard drive development,



improving the reliability, resiliency and overall storage capacity of our drives.

By improving signal-to-noise ratio, IDRC enhances the ability of the hard drive's read channel to distinguish between viable data and random noise. This innovative approach delivers enhanced reliability as well as an 8% increase in drive capacity when compared with previous-generation read channel electronics.

HGST is the first company in the industry to introduce IDRC technology into a production hard drive.

In its first generation, this powerful new technology delivered improvements in signal-to-noise ratio that totaled a full 1 decibel (dB), with second-generation testing yielding a second 1dB increase. In comparison, other signal-to-noise enhancement processes offered increases measured in a few tenths of a dB at best.

The new IDRC technology offers several key benefits, including:

- The most dramatic advance in signal-to-noise enhancement in two decades of hard drive development
- Up to an 8% increase in drive capacity by improving the read channel's ability to distinguish between viable data and random interference
- Enhanced data integrity under adverse operating circumstances and reduced potential for data loss caused by media defects or damage
- Increased reliability through the processing power and improved buffering capabilities of powerful 65nm SOCs
- Full next-generation support for both SATA and SAS-based hard drives

Until now, the design of a device that leveraged IDRC for enhanced reliability was not practical, because the available integrated circuit technologies were unable to deliver the necessary density, power and cost requirements. With the introduction of 65nm mixed-signal CMOS technology, it is now possible to bring a

high-performance, yet cost-effective drive using IDRC technology to market. By leveraging the performance enhancements offered by 65nm processors, including a compact physical form factor, efficient power utilization and extensive availability of onboard RAM, HGST is leading the way to the next generation of high-performance hard drives.

Background

Leadership in signal-processing technology is the corporate DNA of HGST. In 1990, the hard drive division of IBM (now HGST) was the first company to introduce Partial-Response Maximum-Likelihood (PRML) technology in a production hard disk drive. Prior to 1990, all hard disk drives used peak-detection read channels. This technology was relatively easy to understand. A '1-bit' was stored on the disk as a change in magnetism, generating a pulse at the read head output. A '0-bit' was stored as no change in magnetism, generating a flat response at the read head output. The read channel detected the presence of pulses and their locations, enabling the recovery of stored sequences of binary ones and zeros. This scheme required the relative isolation of magnetic pulses from one another, effectively limiting the increase of linear density and the overall storage capacity of the drive.

InterSymbol Interference (ISI) is the term used to describe the overlap of information pulses when magnetic transitions are written too close together. The technology that embraced ISI and provided an optimal detection scheme in the presence of ISI was called PRML. PRML uses equalization circuits to shape the signal from the read head into a crisp mathematical form called a target polynomial, which describes how much each bit overlaps into adjacent bits.


Target polynomials are generally programmable and are chosen so that noise in the final equalized signal is uncorrelated from bit to bit. Under this condition, a Viterbi detector can be used to find the most likely bit sequence that would have generated the read signal through use of a least-squared error metric. The advantage that PRML has over peak-detection technology is that ISI is properly accounted for, allowing linear density to be increased while still maintaining optimal data detection.

There have been incremental enhancements to PRML over the last decade, primarily in the areas of distance-enhancing codes and data-dependent noise. By using simple short block parity where one parity bit is added to 60 consecutive bits of data, a 61-bit codeword can be formed that is different from any other 61-bit codeword in at least two bit locations. This distance-enhancing parity constraint can be incorporated into the Viterbi detector trellis, or included as part of a secondary parity post-processor circuit. The net result is improved performance even after the added redundancy is taken into account.

Data-dependent noise has become more dominant over traditional electronics noise in the read head signal. Transition position jitter is one such example of data-dependent noise, which is generated when the location of a magnetic transition is shifted early or late by the domain boundaries in the disk media. Since this noise is not generated when there are no transitions, the noise is said to be data-dependent. Pattern Dependent Noise Predictive-Partial Response Maximum Likelihood (PDNP-PRML) effectively provides multiple target polynomials that are selected on the data sequence's context for optimal performance.


Achieving Enhanced Performance Through IDRC

IDRC technology includes the same PDNP-PRML foundation used by previous read channel electronics. The new enhancement is an iterative detector that uses a Low Density Parity Check (LDPC) code to correct additional bit errors using successive iterations of the detector. IDRC calculates reliability values (known as soft information) for each data bit read, then tracks and modifies these values based on a message-passing information algorithm determined by the LDPC code.



Soft Information

If ideal binary bits are represented in black and white, then soft information can be thought of as a gray scale that describes detected bits. Soft information is calculated to determine the reliability of each bit, and is usually represented as a Log Likelihood Ratio (LLR). This is the natural logarithm of the ratio, based on the probability that the bit is a 1 divided by the probability that the bit is a 0. The LLR format provides accurate manipulation of soft information using only simple addition, subtraction and compare operations. The LLR format also has desirable dynamic range properties that enable LLR values to be represented in a small number system. The final decision of whether a bit is considered a 1 or a 0 is based on whether the sign of the LLR value is positive or negative. An LLR value of zero indicates the complete inability of the system to determine a bit's binary value.




LLR values can be graphically represented by bars. Bars above the zero reference line are considered BIT=1. Bars below the zero reference line are considered BIT=0. This is illustrated in the graph to the left.

Iterative detection schemes are reliant on accurate soft information. A Soft

Output Viterbi Algorithm (SOVA) is typically employed to generate LLR values for each data bit. Since the Viterbi algorithm is based on sequences and not individual bits, the SOVA block must look both forward and backward from a current bit in order to calculate an accurate LLR bit value.

Coding


As illustrated below, three levels of coding are used in the IDRC data path. This particular ordering of codes is a recent development in the drive industry and is referred to as reverse concatenation. Older hard disk drives used a forward concatenation scheme where the Modulation and Reed-Solomon code blocks are swapped. Reverse concatenation results in a more efficient coding structure and is also an enabling component for the move to iterative detection. Each of these coding levels has an important role in the function of the read channel.



The Modulation code is the outermost level of coding and is very efficient, maximizing format efficiency and adding only 0.5% redundancy to the data stream. Its task is to guarantee that the read signal has continuous information content. This allows both timing and gain information to be extracted from the read signal, making the data stream self-clocking. The Modulation code also prevents stressful patterns from being written to the disk.


The Reed-Solomon code block provides on-the-fly (OTF) correction for errors that remain after the iterative detection process. Up to 16-symbol error correction is guaranteed OTF. Soft information-aided Reed-Solomon decoding is also available in a recovery mode for increased error correction capability. The Reed-Solomon code block also ensures that the Modulation code operates on error-free data.

The binary LDPC Inner code is integral to the iterative detection process. This code relates sets of bits by simple XOR parity. A representative example would be to add 5 LDPC parity bits to 25 bits of data. The resulting 30 bit-long LDPC codeword would consist of 5 sets of 6 bits, where the 6 bits in each set are related to each other by XOR parity. LDPC codes are often described by a parity check matrix as shown below:



The columns represent the bit locations, and the dots in each row show the bits related by parity. In the representative example, each bit (column) is involved in exactly one of the parity checks (rows). This is referred to as a column weight one code. It is also common to have LDPC codes where each bit is involved in two, three or four parity checks. Several properties must be considered in order to achieve an optimum LDPC code. In general, the parity check matrix for an LDPC code looks random, but a deterministic construction is typically employed. This enables the code to be efficiently implemented in hardware. More highly-optimized LDPC codes can be constructed when the block size of the code is above a certain threshold.


Another important aspect of iterative detection is that information is passed via the parity checks of the LDPC code, as well as by the ISI energy coupling between adjacent bits. The LDPC code defines the information transfer paths. A path in the LDPC code that forms a loop (or cycle) will degrade performance by creating a localized form of positive feedback in the iterative process. The diagram below displays several short loops in the code, which are circled in red.



These loops involve information passed via the code's parity checks (horizontal lines), and also by the ISI energy coupling of the adjacent bits (slanted vertical lines). A good code design eliminates all short loops of this type. Ideally the iterative detector will iterate numerous times without generated information flowing back to its point of origin.

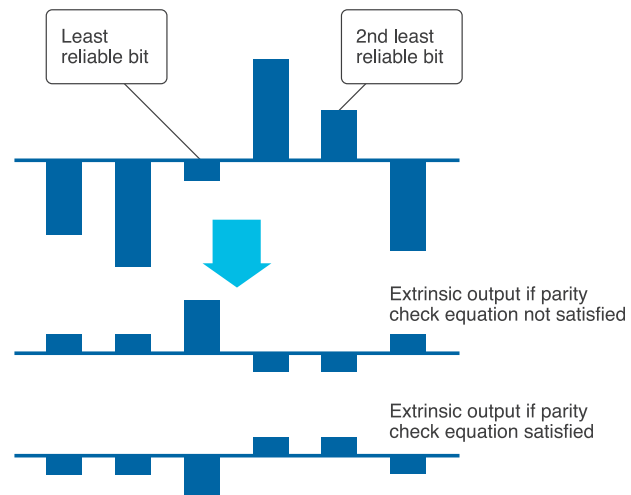
Iterative Detector

As the name implies, the iterative detector performs multiple passes: first through a SOVA block, then followed by a Message Passing (MP) block as shown in the diagram below. The SOVA block generates LLR values for each detected bit. At a given iteration, the sign of these LLR bit values represent the best estimate of the detected binary data sequence. The MP block generates extrinsic LLR information that can be viewed as a delta or bias value, and then used to modify LLR bit values. Each iteration of the detector works to correct additional bit errors, and to reduce the number of unsatisfied parity check equations in the LDPC code.



Both the SOVA and MP blocks pass information between bits. The SOVA block exchanges information between bits that are related to each other by proximity (ISI). The MP block exchanges information between bits that are related to each other by the parity check equations of the LDPC code.

Though the MP algorithm is based on a highly theoretical foundation, it is easy to get an intuitive feel for its operation. The MP algorithm looks at the LLR values for a set of bits related by parity and generates the extrinsic bias values for these same bits, which are then used in the next iteration of the detector. For a simple column weight one binary LDPC code, the MP block operates sequentially on each parity check equation. There are two steps to the operation for the MP algorithm: a bit-to-check step followed by a check-to-bit step. The following diagram helps to illustrate the MP block operation.



It should be noted that a parity-check equation relates a set of scattered bits by parity. The bit-to-check step can be viewed as an information-gathering function. One piece of information gathered is whether or not the parity-check equation is satisfied, and the sign of each LLR bit is used for this purpose. The bit-to-check stage also identifies the least- and second-least reliable bit locations along with their LLR magnitudes.

The check-to-bit step is used to generate the extrinsic output for the MP block. If the parity-check equation is satisfied, the extrinsic information will have a polarity that confirms the current hard decision value for each of the bits. If the parity-check equation is not satisfied, the extrinsic information for each bit will have a polarity that attempts to flip the current hard decision value for each of the bits.

The magnitude of the extrinsic outputs is determined by a simple operation. The extrinsic output for the least-reliable LLR bit has a magnitude proportional to the second-least reliable bit, and the extrinsic output for all other bits will have a magnitude proportional to the least-reliable LLR bits. The iterative application of these simple operations enables the parity check equations to be satisfied and the bit errors to be corrected.

Data Integrity

The robustness of the read channel and the way in which it supports the integrity of the drive's stored data are of primary concern. The first-generation IDRC devices from HGST will continue to use the Reed-Solomon error-correcting code. Reed-Solomon error correction is known for its reliability and powerful error-correction capability, both in terms of random errors and large erasures. IDRC-based solutions have been shown to handle defects and erasures with greater effectiveness when compared with the previous read-channel architecture.

Looking Forward

In the future, IDRC-enabled drives from HGST will utilize more powerful LDPC codes with a greater number of decoding iterations, opening the door to the potential elimination of the Reed-Solomon error correcting code. Based on a single LDPC code instead of the concatenation of two separate codes, this unified coding approach will enable HGST to deliver further enhancements in reliability, resiliency and capacity as hard drive design continues to evolve.



References in this publication to HGST products, programs, or services do not imply that HGST intends to make these available in all countries in which it operates. Some countries have laws that may affect encryption-enabled devices. HGST assumes no responsibility for compliance with laws that apply to our products after they are sold or outside of our control. Accordingly, we suggest that you review all laws of the countries where you conduct business.

Product specifications provided are sample specifications and do not constitute a warranty. Information is true as of the date of publication and is subject to change. Actual specifications for unique part numbers may vary. Please visit the Support section of our website, www.hgst.com/support, for additional information on product specifications. Photographs may show design models.

WPIDRC08EN-02

© 2012 HGST, a Western Digital company

HGST, a Western Digital company
3403 Yerba Buena Road
San Jose, CA 95135 USA

Produced in the United States 11/08.
All rights reserved.