Larger Disk Blocks or Not?

Steve McCarthy
Maxtor Corporation, Longmont, CO 80503
<steve_mccarthy@maxtor.com>

Mike Leis
Maxtor Corporation, Shrewsbury, MA 01545
<mike_leis@maxtor.com>

Steve Byan
Maxtor Corporation, Shrewsbury, MA 01545
<stephen_byan@maxtor.com>

Introduction

The recent annual compound growth rate of disk drive areal density has been 100% - a doubling of capacity every year. This growth rate is faster than Moore's Law - advances in disk technology have been outpacing advances in semiconductor technology. Part of the reason for this spectacular growth rate is that areal density is a two-dimensional problem. Succeeding product generations increase both the number of tracks per inch (TPI) radially and the number of linear bits per inch (BPI) circumferentially. However, both parameters are facing technical challenges that may slow the rate of capacity growth. In this paper, we will briefly examine some of the obstacles to increased BPI and propose an increase in sector size as an aid to surmounting them.

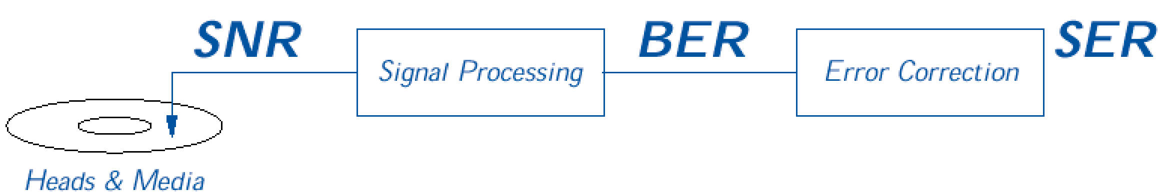

Figure 1: block diagram of heads/media, channel, and ECC

Figure 1 illustrates the communications channel for magnetic recording. There are three areas where improvements can increase the linear bit density: disk heads and magnetic media, signal processing in the read channel, and the error correction coding of the disk block. Improved heads and media could deliver the same signal-to-noise ratio (SNR) at increased BPI. Improved signal processing could deliver the same bit error rate (BER) with worse SNR. Improved error correction could deliver the same sector error rate (SER) with worse BER. In practice, continual improvement in all three areas has helped the recent growth spurt in disk capacity.

Impediments to Linear Bit Density Growth

The disk drive industry faces a challenge to maintain the recent areal density growth rate at a competitive price. Improvements in magnetic recording technology may be slowing. The yields are dropping in heads and media; to avoid price increases the disk drive industry may have to live with lower SNR. The gains in signal processing are slowing and the pitfalls in the new decoding methods currently under research are unknown, so signal processing may not be able to make up for the loss in SNR. As a result the bit-error-rate from the channel may increase.

However, the sector error rate cannot increase. It is the unrecoverable error rate experienced by the customer. It must be maintained or improved. In current drives, the sector error rate is specified to be less than about 10^{-12}.

Allowing the signal-to-noise ratio and the bit error rate to worsen helps continue the growth in areal density. However, error correction would have to improve in order to maintain the sector error rate in the face of worse SNR and BER.

Potential ECC Gains

As shown in figure 2, the gain from error correction has a maximum for 512 byte sectors. For a given set of heads and media, increasing the BPI by 5% reduces the SNR by about 1 dB. The gain in bit error rate from ECC can be translated to an equivalent increase in SNR. The increase in SNR with additional ECC overhead is not linear; the marginal gain decreases with increasing ECC overhead. At some point, the per-sector overhead of the ECC requires an increase in BPI which results a greater SNR loss than the gain from the ECC. Above this crossover point, increasing ECC results in a net decrease in capacity. For 512 byte blocks, this crossover point is at about 54 bytes.

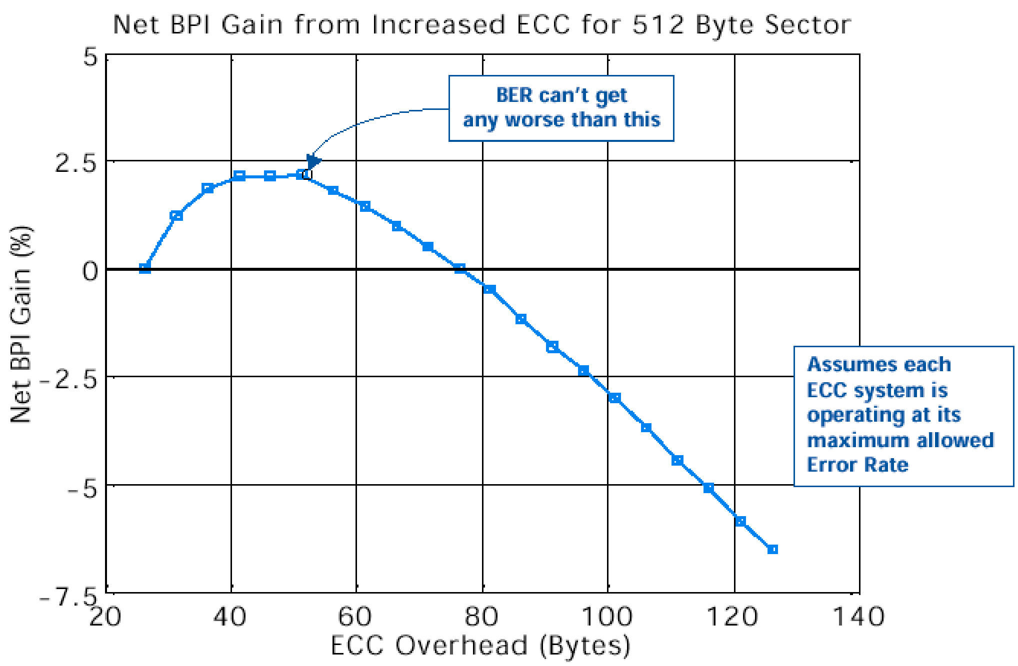

Figure 2: BPI gain vs ECC overhead for 512 byte sectors

Gain from Increased Sector Size

A consequence of information theory is that longer blocks improve the error correction capability for a given ECC overhead. Applying this to a larger disk sector size allows a disk drive operating point with lower signal-to-noise ratio and higher bit error rate. A larger sector size also improves the track format efficiency by amortizing the sector overhead fields over a larger data block. As shown in figure 3, this enables a higher effective bit density with the same heads and media and read-channel.

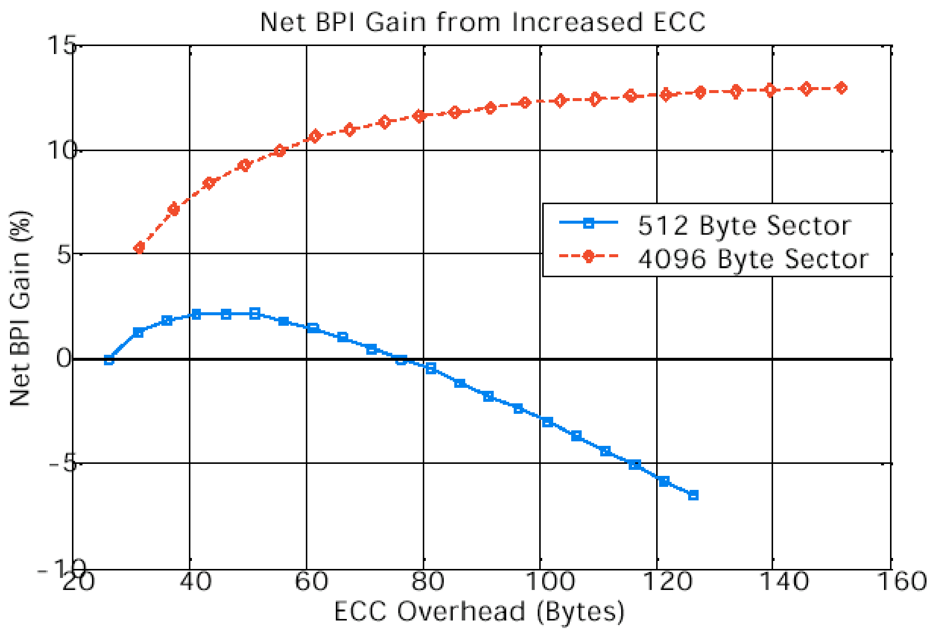

Figure 3: BPI gain vs ECC overhead for 512 and 4K byte sectors

Putting the two together, a 4096 byte disk sector gives a one-time capacity gain of about 10% to 12%. Compared to the current 100% ACGR for disk capacity, this is a small gain. However, it may become more important if the growth rate slows.

There is much controversy over the need to move to larger sectors. Some in the industry believe that 4096 byte blocks are required in order to provide continued capacity growth; we do not share that view.

Summary

Longer sectors give a boost to areal density. The power of an error-correcting code is increased by applying it to a larger block. This improved ECC power allows a higher bit error rate, which allows a lower signal-to-noise ratio, which allows a higher linear bit density for a given set of heads and media and read channel. Longer sectors amortize the per-sector overhead over more data, increasing the track format efficiency. Moving to a 4096 byte sector gives a one-time capacity gain of about 10%. However, some in the industry believe that the the change is required for continued capacity growth.

The International Disk Equipment Manufacturers Association (IDEMA) has proposed moving to 4096 byte sectors in disk drives. This is a user-visible change to disk drives. The atomic unit for writes would be 4096 bytes. Writes of 512 byte blocks would require a costly read-modify-write cycle, either in the drive or in the operating system device driver. Consequently, drive manufacturers would probably continue to offer drives formatted for 512 bytes (with a 10% capacity loss), as well as the 4K byte format.

We know there are significant operating system and software issues with a larger sector size, which we do not present here. We do feel these issues need wider discussion and solicit feedback on the feasibility of moving to a larger sector size. However, we have presented some of the issues motivating the disk drive industry's desire for larger block sizes.