

# HyperLink NAND Flash Architecture for Mass Storage Applications

Roland Schuetz, HakJune Oh, Jin-Ki Kim, Hong-Beom Pyeon, Steven A. Przybylski\* and Peter Gillingham

MOSAID Technologies Incorporated, 11 Hines Road, Ottawa, Ontario, Canada K2K 2X1

\*Verdande Group, Inc. 5630 Meadow Lane, Ann Arbor, MI 48105

## Introduction

The NAND flash memory market is experiencing growing demand from an expanding number of mass storage applications that require high density and high read and write throughput. Prominent among these applications are Flash based Solid State Drives (SSDs), which are expected to replace HDDs in many systems. However, conventional flash interfaces and architectures fail to provide enough bandwidth to keep up with newer serial and Fiber Channel disk-level interfaces. Forecasts indicate that, between now and 2011, the media transfer rate (MTR) for SSDs – the used bandwidth between the controller and conventional flash memories – will grow at less than 5% per year (Fig. 1). In contrast, the interface transfer rate (ITR) – the used bandwidth over the system bus – is expected to grow at 12% per annum over the same period. Barring any breakthrough innovations, SSD performance will increasingly be constrained by the available Flash memory bandwidth (Fig. 1). In addition, the finite number of program-erase cycles offered by Flash memories necessitates the use of wear-leveling algorithms that increase the complexity of SSD implementations and reduce their performance. The HyperLink NAND Flash Architecture facilitates high performance Flash memories for SSD and other applications [1]. It combines a flexible and low pin count interface with an improved core organization and functionality that together deliver high sustained bandwidth and reduced system cost.

## Overall Architecture

The HLNAND Flash Architecture is divided into two independent components: the interface architecture and the memory core architecture (Table 1). The HyperLink Interface Architecture includes a low pin count ring-topology channel with a packet-oriented protocol that is independent of the memory type, organization and functionality. The HLNAND Flash Core Architecture includes significant innovations to both the memory organization and the erase/program functionality. The memory is organized into functionally separate banks, each capable of read, program and erase operations independent of concurrent operations in the other banks. In addition, new erase and program functionality such as page-pair erase, partial block erase and random page program improve system performance by reducing the number of Flash operations needed to implement the critical system-level drive functionality. Additional features such as a novel low-stress program scheme and low Vcc operations improve device longevity.

## Interface Architecture

The Hyperlink physical interface is a high-speed double data rate ring-topology channel with between 1 and 8 data lines and two strobe signals (DSI and CSI). The used link width is set by the controller based on the capabilities of the memories, including the number of data pins each supports. Figure 2 shows a memory system with 3 memories resident on a 4-bit channel. All communication around the ring is point-to-point and command and data packets flow all the way around the ring beginning and terminating at the controller. To save system-level power, command packets can be truncated by a recipient if it is sure that no downstream memories need to see the full contents of the command packet.

The Hyperlink Interface Architecture defines two classes of HL compliant devices. HL1 devices are intended for initial and lower performance applications. It uses LVCMOS signaling and parallel clock distribution to achieve relatively modest operating

frequencies. For example, a dual link, 4-bit HL1 NAND memory system may achieve an MTR of up to 266MBps (DDR), which keeps pace with SATA 3.0 interconnect. HL2 devices use HSTL signaling and source synchronous clocking to achieve much higher operating frequencies. Figure 3 shows an HL2 memory system, and Table 2 lists the primary characteristics of HL1 and HL2 devices. HL1 memory systems will be well suited to the high throughput requirements of SSD and other mass storage applications from the outset; a transition to HL2 interfaces will propel Flash memories into even higher performance regimes that were not previously attainable by Flash memory systems.

An important consequence of the strict bifurcation of the HLNAND Flash Architecture into interface and memory core, is that the HyperLink Interface Architecture is memory-type independent. In fact, the HLNAND Flash Architecture is the first of a series of HL-compatible memory and non-memory architectures that will yield memory and peripheral subsystems consisting of a variety of device types cohabitating on a single link.

## Core Architecture

The HLNAND Flash Architecture defines a wide range of possible densities and organizations. For example, Figure 4 shows the possible implementation of a 16Gb HLNAND device. The device contains two independently operating banks each with its own page buffer. The fully independent bank architecture allows for most operations to be performed concurrently in each of the two banks, greatly improving memory system throughput. This is illustrated in Figure 5 which depicts a ring with four devices carrying out read and program operations on each of the 8 separate banks in order to keep them fully utilized. Table 3 summarizes the features of the 4Gb & 16Gb HLNAND devices that are expected to be lead devices into the marketplace.

Many of the limitations present in conventional NAND flash are addressed by innovations in the memory organization and command set that are part of the HLNAND Flash Core Architecture. For example, a novel low stress program scheme and low Vcc operations minimizes cell Vth shift in unselected pages resulting in higher reliability, greater device longevity and a reduction in power consumption [2]. Support for random page-pair erase, partial block erase, and random page program provides uniform read and write operations and dramatically reduces page copy overhead for wear-leveling and enhances overall system performance.

## Conclusion

The dramatic price reduction of NAND Flash devices in recent years has created an opportunity for Flash to penetrate mass storage applications. This will happen provided the memory vendors can deliver NAND Flash devices with adequate performance and no intrinsic cost premium over the lowest cost conventional NAND Flash devices. The new HLNAND Flash Architecture facilitates this transition by enabling high performance NAND Flash devices with increased longevity and a cost advantage stemming from the low pin count interface and small die size.

## References

- [1] [www.hlnand.com](http://www.hlnand.com)
- [2] Jin-Ki Kim et al., "Low Stress Program with Single Wordline Erase Schemes for NAND Flash Memory", NVSMW, 2007.

