



Phase Change Memory (PCM): A new memory technology to enable new memory usage models

Advancing applications drive memory technology requirements that are met by PCM

Key design requirements for electronic systems

Bandwidth partitioning

PCM scalability

PCM in embedded systems

PCM in wireless systems

PCM in solid state storage subsystems

PCM in computing platforms

Conclusion

Phase Change Memory (PCM) : A new memory technology to enable new memory usage models

Today's advanced applications are driving the demand for new memory technology capabilities. As electronic systems require greater volumes of code and data, the resulting impact creates an insatiable appetite for memory. Phase Change Memory (PCM) meets the needs of today's demanding electronic systems with innovative key technology features.

Key design requirements for electronic systems

Phase Change Memory (PCM) technology directly addresses the needs of today's electronic systems:

Density

With the convergence of consumer, computer and communication electronic systems, an exponential growth of code and even faster growth of data is occurring in all electronic systems. To accommodate this growth, memory density ranges must not only meet current needs but must demonstrate the ability to scale to larger densities as required over time.

Bandwidth and energy

In high-level "convergent" electronic systems, performance is measured in terms of bandwidth, to speed up internet connection, and reduction of power consumption, to enhance mobile use.

The memory system design must support the increasing requirements of bandwidth and reduced power consumption. Non-volatile solid state memory is the proven way to reduce power consumption, as demonstrated with traditional NOR flash memory.

Memory subsystem architecture

Memory subsystem architecture is a key challenge for embedded systems designers. Memory parameters such as density, performance, packaging and interfaces all play a significant role in system-level performance. With the variety of memory types available to the system designer, it is viable to partition the memory subsystem according to the specific needs of the higher-level system and application components. In some cases, caching

is a reasonable approach to achieve an appropriate balance of performance, power and cost. In other cases, partitioning according to the unique characteristics of the underlying memories becomes a more reasonable approach. For example, putting bit-alterable content in a bit-alterable memory instead of attempting to manage bit-alterability in a block-alterable memory.

Bandwidth partitioning

At a high level, there are three main bandwidth categories: Code, data streaming and data storage.

Code: The main variable for code performance is read speed. Code depends on speed of execution utilizing one of the following modes: Execute in Place (XiP) using NOR flash memory to meet high bandwidth and fast random access reads; or, Store and Download (SnD) using NAND flash memory and DRAM for code densities larger than about 1Gb.

Data streaming: The main factor for data streaming performance is programming speed. Data streaming is typically based on DRAM technology, but can be implemented using NAND flash memory and DRAM for densities larger than 4Gb, mainly for density capability and reduced power consumption.

Data storage: The main considerations for data storage performance are density and data retention. However, because density is growing exponentially, latencies between the different parts of the system have a strong impact on the subsystem performance. Data storage usage models typically utilize NAND flash for densities between 4Gb and 100Gb. Given the strong correlation between cycling and data retention, systems that utilize NAND to achieve high write performance are often faced with the difficulty of ensuring adequate data retention to endure long periods of inactivity.

COMPARISON OF HIGH-DENSITY MEMORY TECHNOLOGIES					
Attributes	DRAM	PCM	NAND	MLC NAND	HDD
Non-Volatile	No	Yes	Yes	Yes	Yes
Erase Required	Bit	Bit	Block	Block	Sector
Software	Simple	Simple	Complex	Very Complex	Simple
Power	~W/GB	100→500mW/die	~100mW/die	~100mW/die	~10W
Write Bandwidth	~GB/s	1→100+ MB/s/die	10→100 MB/s/die	~10 MB/s/die	200→400MB/s
Write Latency	~20-50ns	~1μs	~100μs	~800μs	~10ms
Write Energy	~0.1nJ/b	<1nJ/b	0.1-1nJ/b	<1nJ/b	>10nJ/b
Read Latency	50ns	50 - 100 ns	10-25 μs	25-50 μs	~10ms
Read Energy	~0.1nJ/b	<<1nJ/b	<<1nJ/b	<<1nJ/b	>10nJ/b
Idle Power	~W/GB	<<0.1W	<<0.1W	<<0.1W	<10W
Endurance	∞	10 ⁸	10 ⁵ →10 ⁴	10 ⁴ → ?	∞
Data Retention	ms	Not f (cycles)	f (cycles)	f (cycles)	Not f (cycles)

PCM scalability

System designers continue to face significant challenges to design reliable embedded and storage systems based on flash memory. With each new generation, the capabilities of existing memory technologies degrade, requiring significant system-level changes to maintain system-level reliability and performance. Both NOR and NAND flash rely on memory structures which become increasingly difficult to shrink at smaller lithographies. This scaling effect is often referred to as Moore’s Law, where memory densities double with each smaller generation. PCM, however, is based on a physical state change of a chalcogenide material, commonly referred to as GST. Chalcogenide films have already been proven to have stable characteristics to a 5nm node¹. As the PCM memory cell shrinks, the volume of GST material involved in the state change shrinks resulting in reduced power consumption or higher write performance. This unique feature of PCM technology supports the promise of scalability beyond that of other memory technologies.

PCM in embedded systems

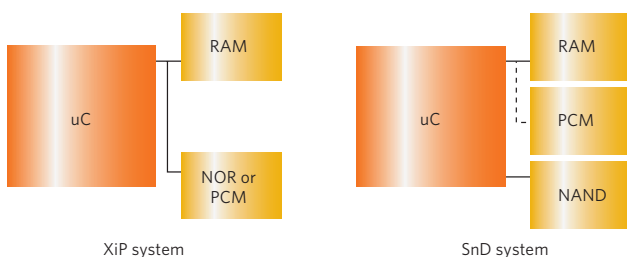


Figure 1. SnD and XiP system architectures

A common use of memory in an embedded system is for code storage. Systems requiring a relatively small amount of memory, less than approximately 2Gb, are architected such that code is executed directly from the NOR flash (XiP). This memory is often

also used as storage memory for an embedded file system. DRAM is often used in these types of systems as a scratchpad memory.

In these types of systems, PCM can be used as a code execution memory. With its bit-alterable feature, PCM is able to displace some or all of the DRAM required in the system. (See Figure 1)

In SnD memory systems, PCM can reduce the density requirements for DRAM while fulfilling the density requirement of the NAND flash. At the same time, the presence of PCM in this type of system simplifies and improves the performance of file systems stored in the PCM due to the bit-alterability and low latency features.

PCM in wireless systems

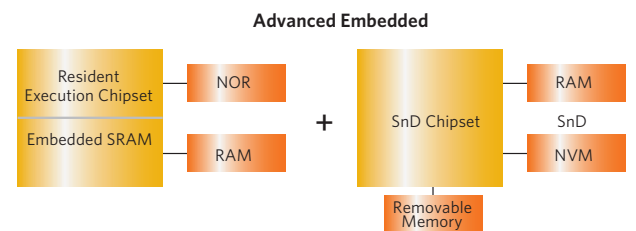


Figure 2. PCM can enhance advanced embedded systems, as found in high-end wireless systems

It is common to find nearly independent subsystems for baseband and application processing in wireless systems. At the highest level, these can be considered as independent embedded systems. Generally speaking, both subsystems have the need for a resident execution memory and for storage of small data structures. In many cases, the applications subsystem is also expected to store and perform operations on larger multimedia content. (See Figure 2)

The low read latencies and fast memory overwrite capabilities of PCM make it an ideal non-volatile XiP solution that scales from low- to high-density wireless solutions. With read latencies that

are slower but on the same order of magnitude as the latencies of DRAM, albeit on smaller page sizes, PCM can serve as an outstanding code execution memory and outstanding read-mostly memory for all but the most frequently manipulated data structures. The bit-alterability of PCM eliminates the need for block erase, which reduces the DRAM requirements even further, resulting in a lower cost memory subsystem.

PCM promises a scalable memory subsystem solution that provides the best overall cost while meeting the increasing performance demands of high-end, multimedia wireless devices.

PCM in solid state storage subsystems

Managing NAND flash in solid state storage (SSD) subsystems is a challenge due to the block-alterable nature of the NAND technology. It is also challenging to handle increasing levels of error management required when the memory is heavily program/erase cycled or frequently read.

PCM can be used in SSD systems to store frequently accessed pages and to store those elements which are more easily managed when manipulated in place. Examples of these types of elements include: parity bits for data stored in NAND, bad block tables, and block and page mapping tables. In this scenario, a small amount of PCM could be used to enhance the manageability of NAND. (See Figure 3) By minimizing the stress on the NAND memory, higher density MLC NAND is enabled, thus leveraging the capability of PCM to lower the cost of the NAND flash in the subsystem. This caching with PCM will improve the performance and reliability of the subsystem.

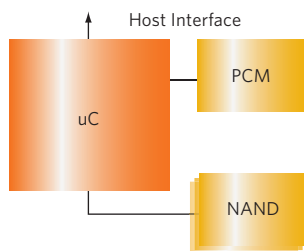


Figure 3. Hybrid solid state storage

Additionally, when erased pages are scattered across many blocks (near full state), PCM can provide further reliability improvement. Managing a block-alterable memory in a near-full state implies

that multiple erase cycles are likely required to free space to store the new data being written to the device. This increases the number of cycles on the device and further accelerates the time until the maximum endurance limits are reached.

The bit-alterable nature of PCM solves the issue of increased write cycles when the device is full. Higher endurance of PCM addresses the needs of these systems when heavy use is expected.

PCM in computing platforms

As a volatile memory, DRAM consumes power to simply maintain the contents of the memory. As a non-volatile memory, PCM banks can be turned off when they are not in use to provide reduced power in idle states. More importantly, turning off the banks decouples the relationship between density and power consumption. This results in a PCM subsystem density envelope that is not limited by the power envelope constraints of that system. In addition to non-volatility, PCM offers endurance and write latencies that are compelling for this type of application. This is a key advantage over read-mostly solutions that have been attempted until now.

Conclusion

Phase change memory can be exploited by the memory system and by the convergence of consumer, computer and communication electronic systems. The caching of the existing memory technologies, reducing the overall system cost and system complexity will be the compelling motivation for PCM adoption. Bandwidth will drive the sustaining side of PCM in code and data transfer applications, while reduction in power dissipation will represent a further added value of this technology.

PCM is today's memory breakthrough. Like flash, PCM is a non-volatile memory that can store bits even without a power supply. But unlike flash, data can be written to cells much faster, at rates comparable to the dynamic and static random-access memory (DRAM and SRAM) used in all computers and cell phones today. Quite simply, PCM blends together the best attributes of NOR flash, NAND flash, EEPROM and RAM—delivering a new category of memory for new usage models.

For more information on PCM, please visit Numonyx.com or contact your local Numonyx sales representative.

This paper was authored by Sean Eilert, Numonyx Director of Architecture, Principal Engineer; Mark Leinwander, Numonyx Systems Manager, Principal Engineer; and Giuseppe Crisenza, Numonyx Vice President of Strategic Alliances.

¹PCM scalability as referenced by C. D. Wright, M. M. Aziz, M. Armand, S. Senkader, and W. Yu, "Can We Reach Tbit/sq.in Storage Densities with Phase-Change Media", EPCOS 2004, and C. Lam, SRC NVM Forum 2004.

numonyx.com

Copyright © 2009 Numonyx B.V. All rights reserved. Numonyx and the Numonyx logo are trademarks or registered trademarks of Numonyx B.V. or its subsidiaries in other countries. *Other names and brands may be claimed as the property of others.

0609/NH/BMG/PDF 109133-001 Please recycle ♻️



numonyx™