AMD MirrorBit™

White Paper

July 2003

The following document refers to Spansion memory products that are now offered by both Advanced Micro Devices and Fujitsu. Although the document is marked with the name of the company that originally developed the specification, these products will be offered to customers of both AMD and Fujitsu.

Continuity of Specifications

There is no change to this document as a result of offering the device as a Spansion product. Any changes that have been made are the result of normal documentation improvements and are noted in the document revision summary, where supported. Future routine revisions will occur when appropriate, and changes will be noted in a revision summary.

Continuity of Ordering Part Numbers

AMD and Fujitsu continue to support existing part numbers beginning with "Am" and "MBM". To order these products, please use only the Ordering Part Numbers listed in this document.

For More Information

Please contact your local AMD or Fujitsu sales office for additional information about Spansion memory solutions.

Publication Number **25260** Revision **C** Amendment **0** Issue Date **September 18, 2002**

AMD

MirrorBit™

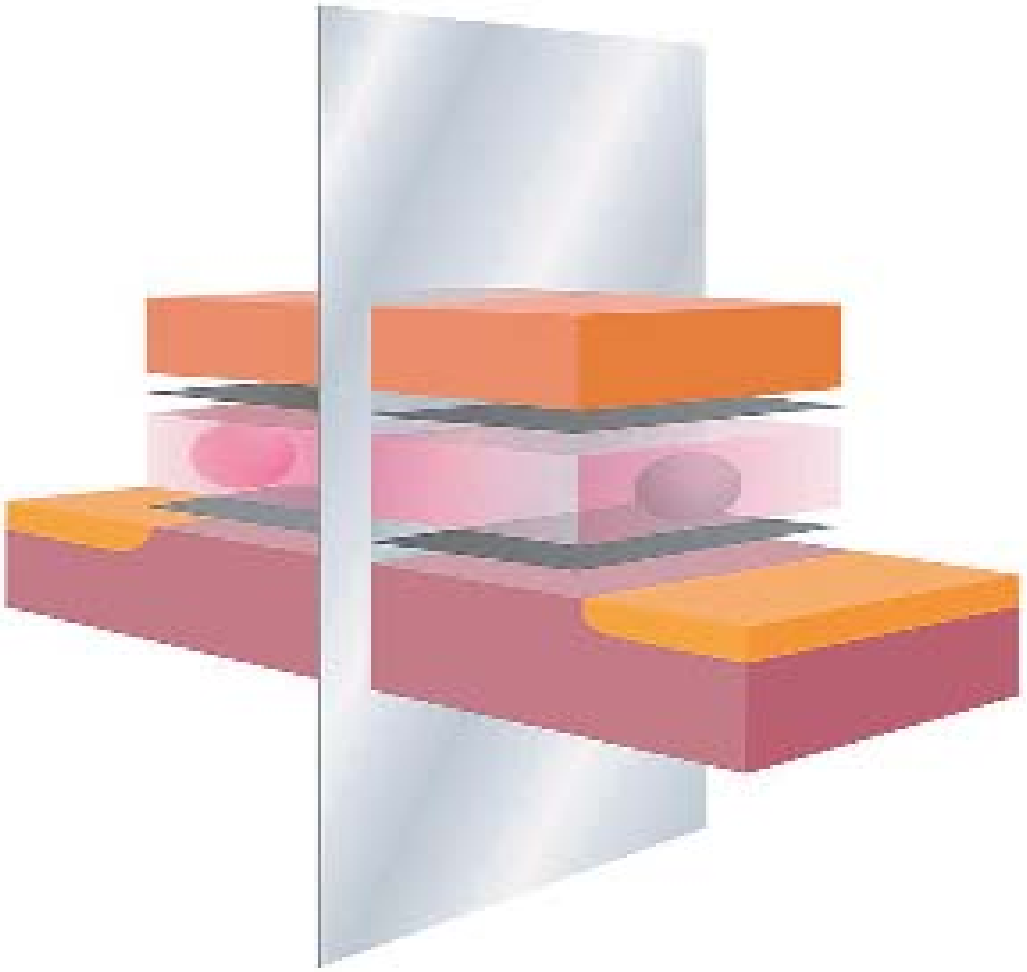

White Paper

AMD FLASH MEMORY

AMD MirrorBit™

White Paper

OVERVIEW

AMD continues to revolutionize the world of Flash memory. In 1991 AMD introduced negative gate erase, a revolutionary architecture that enabled highly reliable, single power Flash memory. In 2001, AMD announced the MirrorBit™ architecture, which – for the first time – enables a Flash memory product to **hold twice as much data as standard Flash, *without compromising device endurance, performance or reliability***. This new architecture offers customers a highly reliable and low-cost, Flash memory solution.

The MirrorBit cell doubles the storage capability of the basic Flash memory cell and therefore enables low-cost, high-density Flash memory products. Unlike competing multi-level cell (MLC) technology, the MirrorBit architecture delivers this enhanced storage without compromising device performance or reliability. This revolutionary architecture is the result of years of research and development on the design, processing, testing and characterization of multi-bit cells that have culminated in AMD's patented MirrorBit architecture.

The MirrorBit architecture's ability to store two bits of data in one cell, without compromising data integrity, is achieved by storing each of the two units of charge individually, and in a different location, within a single memory cell. Since each bit in the MirrorBit memory cell is in a physically different location, the bits are independent and this allows AMD to offer the same performance and reliability as standard single-bit Flash memory products.

Customers can benefit from this revolutionary architecture in several ways:

- Low-cost Flash memory devices
- High-density Flash memory devices
- High levels of reliability
- Performance comparable to today's standard NOR Flash memory
- AMD's award winning service and support

Publication # 25260 Rev: C Amendment/0
Issue Date: September 18, 2002

MirrorBit Cell Diagram

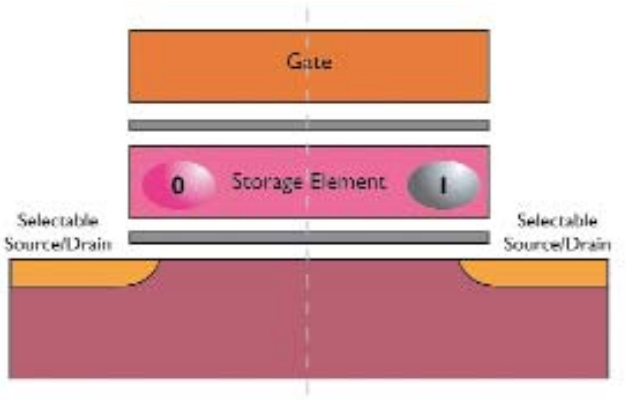

Distinctive Features

- Two bits of data stored in one memory cell. – *This reduces costs and enables low-cost, high-density Flash memory.*
- Simple and symmetric cell architecture. – *This makes the device easy to manufacture.*
- Two distinct charge storage areas within one cell. – *This increases reliability and data integrity.*
- Each unit of charge individually programmed and erased. – *Allows high performance and high reliability.*
- Full charge stored at each end of the cell. – *Enables higher levels of reliability and performance.*
- Products designed to be completely pin-compatible with AMD's current 64Mb LV Flash. – *Provides a seamless migration path from AMD's single bit Flash devices.*

BENEFITS

The AMD MirrorBit architecture breaks the traditional trade-off between cost and performance. For the first time in history, Flash memory designers and users can expect high levels of performance and reliability as well as the industry's leading cost structure. Customers therefore are no longer forced to sacrifice reliability or performance for the sake of cost.

AMD's MirrorBit based products have the same performance, endurance and reliability as AMD's standard 3 Volt products. Products based on this new architecture are completely pin-compatible with AMD's standard 3 Volt (LV product family) and offer a simple migration path for customers from the current LV family of products to higher densities such as 256 Mb and eventually IGb. Since these products are offered in ***the same packages and with the same pinouts*** as AMD's current LV family of products, AMD's customers can easily migrate their existing designs to MirrorBit based Flash.

Most importantly, MirrorBit Flash devices are **available today**, so designers can immediately enjoy MirrorBit technology's many benefits.

In addition to offering low cost and high reliability, products based on the MirrorBit architecture also offer two new performance-enhancing features, a read and write page buffer that increases data throughput. As device densities increase data input and output become a major bottleneck and this limits how fast data gets in and out of the device. Since the MirrorBit architecture is well suited for ultra-high density Flash memory devices, products based on MirrorBit architecture offer these page buffers to further improve functionality and device performance.

AVAILABILITY AND ROADMAP

The MirrorBit architecture has been developed and tested by AMD on a number of different products and densities. The architecture and process technology have met our internal qualification goals. We are currently in the midst of rolling out an entire family of products based on this architecture.

The first 64 Mb, MirrorBit product was introduced in early 2002 and a 256 Mb MirrorBit product is slated for introduction in Q3, 2002. The MirrorBit architecture based product family includes Flash devices from 16 to 256 Megabits and beyond, and is available in a variety of packages such as TSOP, FBGA, and AMD's new Fortified BGA package.

MIRRORBIT TECHNOLOGY

The MirrorBit architecture doubles device density while simultaneously overcoming the fundamental limitations of competing MLC technologies - namely low levels of reliability and performance. A breakthrough design and a relentless commitment to quality and performance solve the classical paradox. The breakthrough itself is achieved by storing electrons in two physically distinct and independent locations within a single memory cell.

The MirrorBit architecture uses a radical new design and process technology to double the density without any performance penalty. In the MirrorBit cell, the basic transistor itself is very different from conventional transistors. Instead of the classical asymmetric transistor with a distinctly doped source and drain, the MirrorBit cell uses a symmetric transistor with similar source and drain. The charge storage element has been modified to allow electrons to be stored on either side of the cell. Once electrons are placed into one side of the storage element, say on the left side of the cell, they remain trapped on that side. Consequently, read, program and erase operations are performed at full speed and power regardless of whether one is using the left bit or right bit. As a result, the basic memory cell behaves as though it were two independent conventional memory cells. This architecture therefore offers twice the density of standard Flash without sacrificing performance or reliability.

The MirrorBit architecture therefore allows AMD to offer the same performance and reliability as standard single-bit Flash memory products. AMD’s customers therefore have access to low-cost, high-density Flash memory solutions without compromising either performance or reliability.

AMD MirrorBit™ Cell

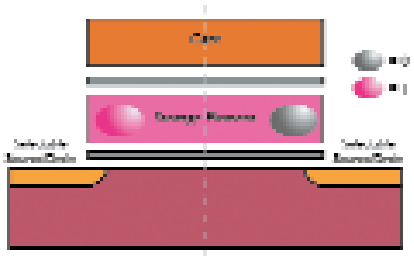

MirrorBit

- 2 bits per cell
- Low-cost structure
- Each bit is ***independent***
- Fast programming
- Fast read
- Easily scalable

Competing Multi Level (MLC) Flash Memory Cell

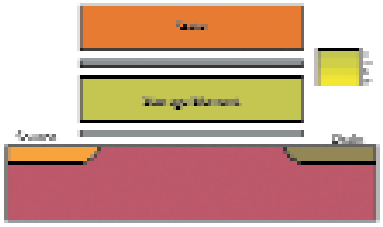

Competing MLC

- 2 levels per cell
- Low-cost structure
- Two bits are ***interdependent***
- Slow programming
- Slow read
- Difficult to scale

Basic Flash Memory Cell

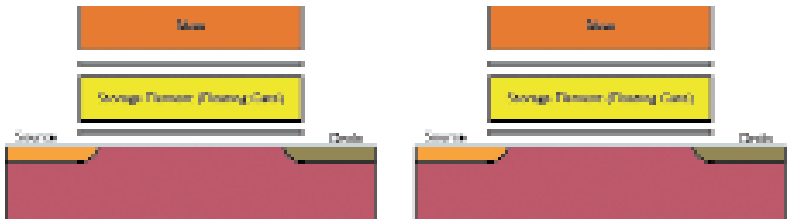

Conventional Flash

- 1 level per cell
- Consumes more silicon
- Fast programming
- Fast read
- Scales with lithography

DEVICE TECHNOLOGY OVERVIEW

In order to fully understand how the MirrorBit architecture works, one needs to understand how basic Flash memory cells work. The following section gives a detailed overview of how Flash memory cells, multi-level cell technologies and the MirrorBit architecture works.

Conventional Flash memory cells are based on MOS (Metal Oxide Semiconductor) transistors, which are probably the most important component of semiconductor devices. This transistor typically consists of a thin, high-quality gate oxide sandwiched between a conducting polysilicon gate and a crystalline silicon semiconductor substrate. The substrate under the gate is lightly doped and is a P-type (electron-deficient) substrate. The source and drain areas are more heavily doped and are N-type (electron-rich).

The floating gate memory cell is very similar to the basic MOS transistor. The difference is in the presence of an additional polysilicon gate. This so-called "floating gate" lies between the control gate and the silicon substrate. Insulating materials such as silicon dioxide surround the floating gate. Flash memory devices store data by holding charge within the floating gate.

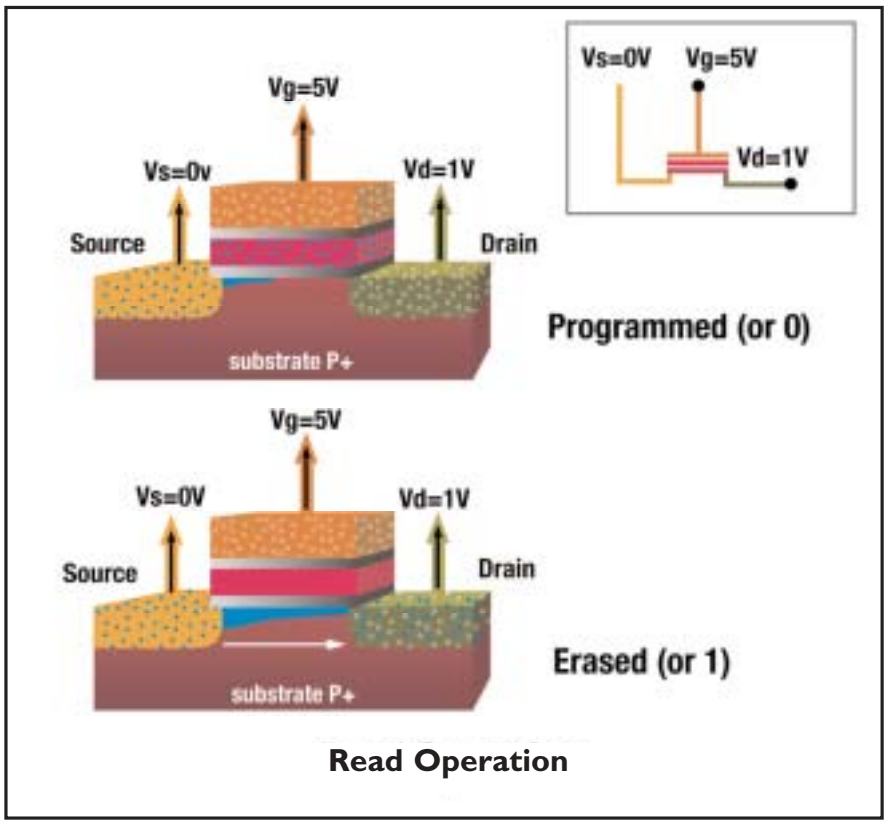

Read Operation

The basic transistor works as follows. In a neutral state, there is no conductive path between the source and drain regions. When a positive voltage is applied to both the gate and drain, a channel begins to form between the source and drain. When the gate voltage is sufficiently large the channel is completely formed and electrons flow from the source to the drain.

The voltage at which the channel forms is called the threshold voltage of the transistor, and is often referred to as Vt. Increasing the control gate voltage beyond the threshold voltage results in an increase in the current flowing between the source and drain.

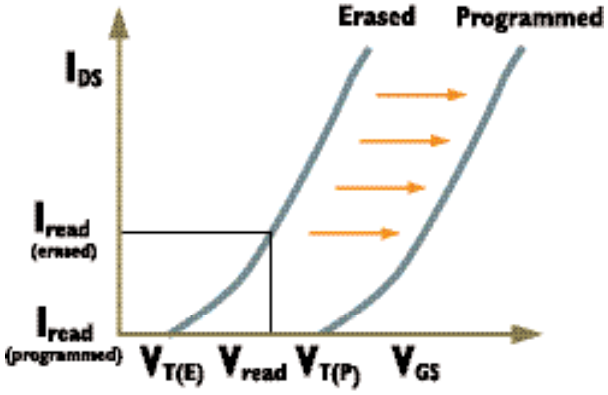

Threshold Voltage

The key to Flash memory cell operation is that when charge is stored in the floating gate, the threshold voltage of the transistor increases. This change in threshold voltage enables one to detect whether or not a cell is programmed.

Single-bit cells basically have two threshold states, a programmed state (threshold voltage V_{tp}) and an erased state (threshold voltage V_{te}). The read voltage is selected to be between the programmed and erased threshold voltages. If current flows at the read voltage, then the cell is read as erased.

COMPETING MULTI-LEVEL CELL TECHNOLOGIES

One common way of increasing the density of memory cells is to use competing multi-level cell (MLC) technology. Competing multi-level cells basically store fractional levels of charge within a cell to offer increased data storage capability. Since the threshold voltage of a transistor depends on the number of electrons in the floating gate, it is possible to store different units of data based on differences in the number of electrons stored in the cell. As a result, multi level cells store more than the basic 0 and I, and instead they store 00, 01, 10 and 11. To store n bits per cell, the device needs to have 2^n different states. Each of these states requires a different number of electrons to be stored in the floating gate, and consequently each state has its own threshold voltage.

Competing MLC Performance:
The read operation for competing multi level cell technologies is intrinsically slow since the cell needs to distinguish between these multiple possible threshold voltages. In order to program a competing multi level cell, electrons must be very precisely metered into the floating gate. If too many or too few electrons are allowed into the floating gate, the data stored in the memory cell is corrupted. Consequently, the program operation is performed very carefully and is inherently slow.

Competing MLC Reliability:
The reliability of competing MLC technology is even more problematic than the reduced program, erase and read performance. Since competing MLC technology depends on differences between small numbers of electrons in the floating gate, a loss of few electrons can cause data corruption and device failure. Competing MLC devices use higher operating voltages to ensure that the basic cell has a sufficiently wide Vt window. However, the associated electric fields cause oxide breakdown and other stresses and consequently limit the endurance of the device.

MIRRORBIT CELL TECHNOLOGY

The MirrorBit architecture doubles device density while simultaneously overcoming the fundamental limitations of competing MLC technologies - namely low levels of reliability and performance. This breakthrough is the result of fundamental changes in device architecture and process technology. These two advances allow the MirrorBit memory cell to store electrons in two distinct and independent locations within a single cell and thereby double device density without compromising performance or reliability.

In the MirrorBit cell the basic unit of charge storage - the memory cell transistor, is different from the conventional floating gate memory transistor. Unlike other Flash memory cells, which use an asymmetric transistor with a distinctly doped source and drain, the MirrorBit cell uses a symmetric transistor with similar source and drain. The floating gate or charge storage element has been fundamentally changed to allow electrons to be stored on either side of the cell. Once electrons are placed into one side of the storage element, say on the left side of the cell, they remain trapped on that side.

Consequently, read, program and erase operations are performed at full speed and power regardless of whether one is using the left bit or right bit. Each bit within the MirrorBit cell has full operating margin and as a result, the basic memory cell behaves as though it were two independent conventional memory cells. This architecture therefore offers twice the density of standard Flash without sacrificing performance or reliability.

MirrorBit Performance and Reliability
The MirrorBit architecture enables one memory cell to hold two virtual transistors. Each of these transistors can be read or programmed independent of the other, and at full power. The threshold voltage margins of the device are also large, and about the same as that of single bit Flash. The performance and the reliability of the device is therefore comparable to the performance and reliability of today's conventional single bit NOR Flash memory. As a result, users can continue to expect the high levels of performance and reliability that they are currently used to.

CONCLUSION

AMD's MirrorBit cell doubles the storage capability of the basic Flash memory cell and therefore enables low-cost, high-density Flash memory products. The MirrorBit architecture delivers this enhanced storage without any compromises. A breakthrough design and a radical new process technology, coupled with a relentless commitment to quality and performance have enabled AMD to create the MirrorBit architecture, which is the result of years of research and development on the design, processing, testing and characterization of multi-bit cells.

The breakthrough itself is achieved by storing electrons in two physically distinct and independent locations within a single memory cell. The MirrorBit architecture enables one memory cell to hold two virtual transistors. Each of these transistors can be read or programmed independent of the other, and at full power. As a result, the MirrorBit architecture delivers the same performance and reliability as conventional Flash memory, but at half the cost.

AMD's MirrorBit products based on this new architecture are completely pin-compatible with AMD's standard 3 Volt (LV product family). These products are offered in ***the same packages and with the same pinouts*** as AMD's current LV family of products, and provide an easy migration path to high-density devices such as 256 Mb and eventually 1Gb. Most importantly, MirrorBit Flash devices are **available today**, so designers can immediately enjoy MirrorBit technology's many benefits.