Please note that Cypress is an Infineon Technologies Company.

The document following this cover page is marked as “Cypress” document as this is the company that originally developed the product. Please note that Infineon will continue to offer the product to new and existing customers as part of the Infineon product portfolio.

Continuity of document content

The fact that Infineon offers the following product as part of the Infineon product portfolio does not lead to any changes to this document. Future revisions will occur when appropriate, and any changes will be set out on the document history page.

Continuity of ordering part numbers

Infineon continues to support existing part numbers. Please continue to use the ordering part numbers listed in the datasheet for ordering.

WHITE PAPER

Ravi Prakash,
Applications Engineer Sr. Principal
Cypress Semiconductor Corp.

Nonvolatile SRAM (nvSRAM) Basics

Abstract

This white paper describes the basic operations of a nonvolatile static random access memory (nvSRAM), using a parallel nvSRAM as an example. The same basic operations also apply to the Cypress serial nvSRAMs.

Introduction

Static random access memory (SRAM) loses its content when powered down, and is classified as volatile memory. The memory is volatile because there is no data when power is restored to the device. Another example of volatile memory is the dynamic random access memory (DRAM) used in all desktop computers and laptops.

Memory that retains its data without power is classified as nonvolatile memory. Examples of nonvolatile memory are nonvolatile SRAM (nvSRAM), ferroelectric RAM (F-RAM™), electrically erasable programmable ROM (EEPROM), and flash memories. This class of memory is used in applications in which critical data must be stored after power is removed or when power is interrupted during operation. An example of power interruptions is the hot plugging of cards in servers, industrial computers, and medical equipment. Several submodules can be unplugged while the equipment is operating (that is, their power can be interrupted) and new submodules can be plugged in (hot plugged because power is present in the equipment) without losing any critical data and/or operations to the equipment.

The nvSRAM is a class of nonvolatile memory that combines SRAM features with nonvolatility. Cypress nvSRAMs do not use a battery or other energy sources to retain data. The nvSRAM has several advantages for applications in which a high speed and nonvolatile storage are required at a low cost compared to alternative solutions that include large supercapacitors and batteries to retain data on devices when power is interrupted. These applications include smart meters, servers, programmable logic controllers (PLCs), gaming, multifunction printers, and storage units.

Cypress offers several families of high-speed, high-performance nvSRAM products that combine the performance characteristics of high-speed SRAM with reliable nonvolatile elements (providing more than 20 years of data retention). The data is retained in the nonvolatile elements that are integrated with each SRAM cell. While operating as high-speed SRAM (< 20 ns access time), the nvSRAM can store data to (STORE) or recall data from (RECALL) nonvolatile elements. The STORE and RECALL operations can be done either by a user command (firmware) or by the device automatically without a user command – AutoStore at power-down and power-up RECALL at power-up.

In the AutoStore mode, the nvSRAM device transfers the SRAM cell data automatically to the nonvolatile elements during power interruptions using the charge from an external capacitor. It then restores the data from the nonvolatile elements into the SRAM cells, without any software intervention, when power is restored to the unit.

The nonvolatile elements (also known as shadow EEPROM) guarantee data retention for a minimum of 20 years at the maximum operating temperature.

nvSRAM Cell Architecture

Cypress nvSRAM combines the standard SRAM cells with EEPROM cells in SONOS technology to provide a fast read/write access and 20 years of data retention without power. The SRAM cells are paired one-to-one with EEPROM cells.

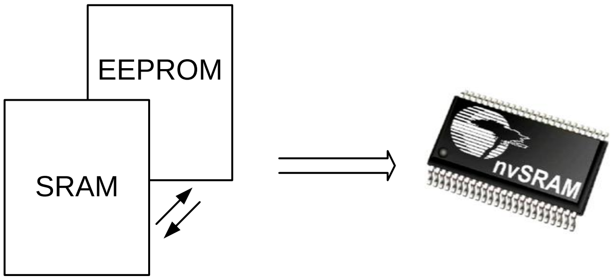

Figure 1. nvSRAM Architecture

The nvSRAMs are in the CMOS process, with the EEPROM cells having a SONOS stack to provide nonvolatile store.

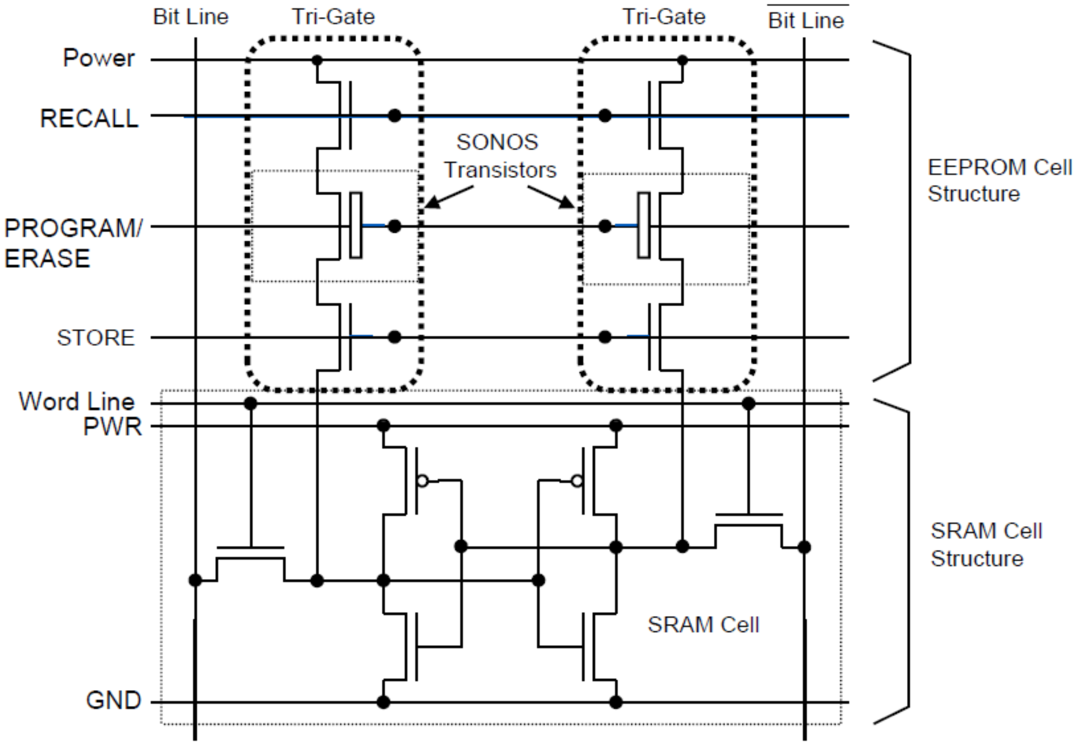

Figure 2. nvSRAM Cell Architecture

Cypress nvSRAM products feature a SONOS nonvolatile cell, which is built on a standard SRAM cell. When normal power is applied, the device looks and behaves in a similar manner as a standard SRAM. However, when power drops out, each cell’s contents can be stored automatically in the nonvolatile element positioned above the SRAM cell. This nonvolatile element uses standard CMOS process technology to obtain the high performance of standard SRAMs. In addition, the SONOS technology is highly reliable and supports 1 million STORE operations.

Device Interface

The interface to the nvSRAM is identical to that of the high-speed SRAM, except for a few extra pins exclusive to the device. Figure 3 shows the logic block diagram of a parallel interface nvSRAM. The address lines, data lines, and control lines ($\overline{OE}$, $\overline{CE}$, $\overline{WE}$) provide the same interface as a high-speed SRAM. For normal read and write operations, the nvSRAM is accessed in the same way as an SRAM. Serial nvSRAMs and asynchronous NAND interface nvSRAMs have the same logic blocks as SRAMs, except that they have a serial interface and NAND interface, respectively instead of a parallel interface.

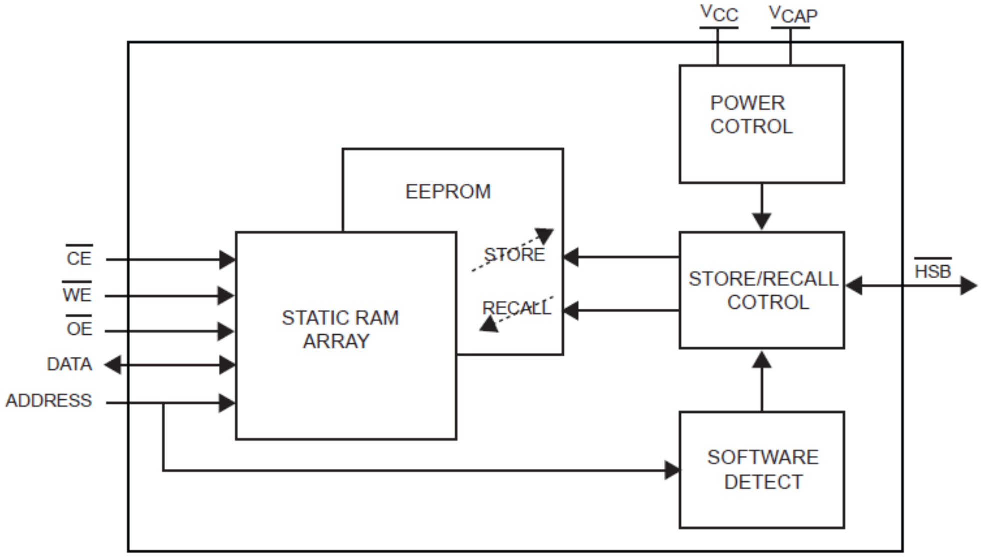

Figure 3. Logic Block Diagram

The power control block detects variations on the power supply (V_{CC}) to trigger an AutoStore operation. The software detect block detects address read sequences for STORE, RECALL, and enable/disable AutoStore operations. The store/recall control block initiates STORE or RECALL operations, including a hardware STORE operation using the $\overline{HSB}$ pin. Software sequences required to initiate nonvolatile operations use standard SRAM control pins. As a result, very few hardware modifications are required to use nvSRAMs instead of standard SRAMs.

The only external component required for nvSRAM AutoStore operations is the capacitor connected to the V_{CAP} pin, which is charged to the supply voltage on power-up. The charge from this capacitor is used to perform an AutoStore operation (transferring the contents of SRAM to nonvolatile elements on power-down). Figure 4 shows the connection of the storage capacitor (V_{CAP}) and the recommended pull-up on the $\overline{WE}$ pin. Because controllers take longer to boot up compared to the 30 ms for the nvSRAM, controller I/Os can be floating when the nvSRAM is ready for read/write. The pull-up on the $\overline{WE}$ pin prevents any inadvertent writes to the parallel nvSRAM due to the floating controller I/Os.

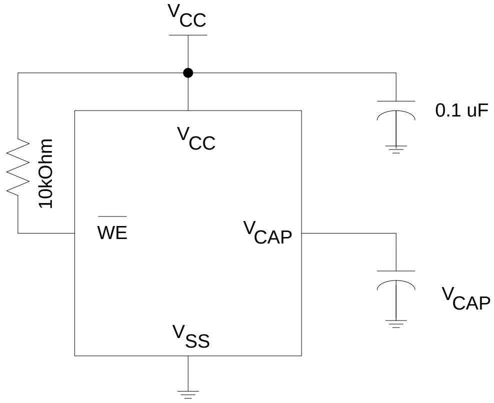

Figure 4. AutoStore Mode

Nonvolatile STORE Operation

A STORE operation is used to transfer the data in parallel from the SRAM to the nonvolatile SONOS cells. For example, in a 4-Mb nvSRAM, the data of all the 4-Mb SRAM cells is stored into the 4-Mb EEPROM elements simultaneously. This parallel transfer of data enables the entire SRAM array to be stored in a maximum of 8 ms. The nvSRAM STORE operation can be initiated in three ways: AutoStore, activated on device power-down; software STORE, activated by a software read sequence; and hardware STORE, activated by the $\overline{\text{HSB}}$ pin.

During the STORE cycle, the previous nonvolatile data is first erased; then the nonvolatile elements are programmed. After a STORE cycle is initiated, further input and output are disabled until the cycle is completed.

The $\overline{\text{HSB}}$ signal can be monitored by the system to detect if a STORE cycle is in progress. The busy status of nvSRAM is indicated by the $\overline{\text{HSB}}$ pin being pulled LOW. To reduce unnecessary nonvolatile stores, AutoStore and hardware STORE operations are internally ignored unless at least one write operation has taken place since the most recent STORE or RECALL cycle. However, software-initiated STORE cycles are performed regardless of whether a write operation has taken place.

AutoStore Operation

In AutoStore operation, data is automatically stored into the EEPROM elements when the system power supply V_{CC} drops below the V_{SWITCH} level.

During normal operation, the device draws current from the V_{CC} to charge the capacitor connected to the V_{CAP} pin. This stored charge is used by the chip to perform a single STORE operation. If the voltage on the V_{CC} pin drops below an internally set level, V_{SWITCH}, the part automatically disconnects the V_{CAP} pin from V_{CC} and a STORE operation is initiated with power provided by the V_{CAP} capacitor. (**Note**: For details about nvSRAM features, refer to the appropriate nvSRAM datasheet. Table 1 provides a quick reference to the major nvSRAM parameters.)

Hardware STORE

The $\overline{\text{HSB}}$ pin is used to initiate a hardware STORE cycle. When the $\overline{\text{HSB}}$ pin is driven LOW by the external controller, the nvSRAM conditionally initiates a STORE operation after a small delay, t_{DELAY}. However, a STORE cycle begins only if a write to the SRAM has taken place since the last STORE or RECALL cycle. The t_{DELAY} time allows for the completion of any write operation in progress. The $\overline{\text{HSB}}$ pin also acts as an open-drain driver (internal 100-kΩ weak pull-up resistor) that is internally driven LOW to indicate a busy condition, when the hardware or software STORE is in progress.

During any STORE operation, regardless of how it is initiated, the nvSRAM continues to drive the $\overline{\text{HSB}}$ pin LOW, releasing it only when the STORE is complete. When the STORE operation is complete, the nvSRAM remains disabled until the $\overline{\text{HSB}}$ pin returns HIGH. The $\overline{\text{HSB}}$ pin can be left unconnected if it is not used.

Software STORE

Data is transferred from the SRAM to the nonvolatile memory by a software address sequence. The software STORE cycle is initiated by executing sequential $\overline{\text{CE}}$-or $\overline{\text{OE}}$-controlled read cycles from six specific address locations in an exact order. During the STORE cycle, an erase of the previous nonvolatile data is first performed, followed by a program of the nonvolatile elements. After a STORE cycle is initiated, further input and output are disabled until the cycle is completed.

Nonvolatile RECALL Operation

A RECALL operation copies data from EEPROM to SRAM, in parallel, within a maximum of 30 ms. The RECALL cycle can be initiated using either of the following methods:

Power-Up RECALL

During power-up or after any low-power condition ($V_{CC} < V_{SWITCH}$), a RECALL cycle is initiated when V_{CC} exceeds the V_{SWITCH} and takes $t_{HRECALL}$ to complete. During this time, the $\overline{\text{HSB}}$ pin is driven LOW by the $\overline{\text{HSB}}$ driver and all reads and writes to nvSRAM are inhibited.

Software RECALL

A software RECALL operation is initiated in a similar manner as a software STORE operation. This is done by reading a specific sequence of six address locations, with no access to other locations in between.

nvSRAM During Nonvolatile Operations

During STORE and RECALL operations, the nvSRAM is not available to the system. All levels and transitions on the input pins are ignored and all data output pins (except the $\overline{\text{HSB}}$ pin) are tristated. After the nonvolatile cycle (software STORE/software RECALL) is completed, a read or write cycle can be initiated immediately.

The $\overline{\text{HSB}}$ pin indicates when a nonvolatile STORE is in progress. This pin is internally driven LOW whenever a nonvolatile STORE is in progress.

If the $\overline{\text{HSB}}$ pin is connected to any other device in the system, a pull-up resistor to V_{CC} must be used on the $\overline{\text{HSB}}$ line. A weak internal pull-up resistor keeps this pin HIGH, if an external pull-up is not connected (optional). The value of the pull-up resistor must be selected so it does not overpower the internal pull-down driver on $\overline{\text{HSB}}$. Typically, a 10-kΩ pull-up is sufficient.

Protection Against Inadvertent STORE Operations

Cypress nvSRAMs have several built-in measures to prevent inadvertent STORE operations. When V_{CC} is below $V_{SWITCH,}$ the nvSRAM disables its I/Os and prevents user-initiated STORE operations. This ensures that STORE operations initiated by $\overline{\text{HSB}}$ and software sequences are not started if the V_{CC} is too low to allow a successful completion.

Hardware STORE and AutoStore operations require that at least one WRITE operation take place since the last nonvolatile operation. This feature ensures that a fluctuation in the supply line or noise on the $\overline{\text{HSB}}$ line does not initiate an unnecessary STORE operation.

nvSRAM Applications

Cypress nvSRAMs are ideal for all applications that require fast access, high reliability, and unlimited endurance. These applications include write journaling in RAID storage systems, smart meters, programmable logic controllers (PLCs), gaming, and multifunction printers (MFPs). The following are two use cases.

Write Journaling in RAID Systems

Figure 5 depicts a generic RAID storage system architecture. The servers connected on the network send the data to a switch, which directs it to different storage boxes. The processor in the storage box also interacts with multiple storage boxes through the switch to enable virtualization. The data received by the processor is sent to the RAID controller.

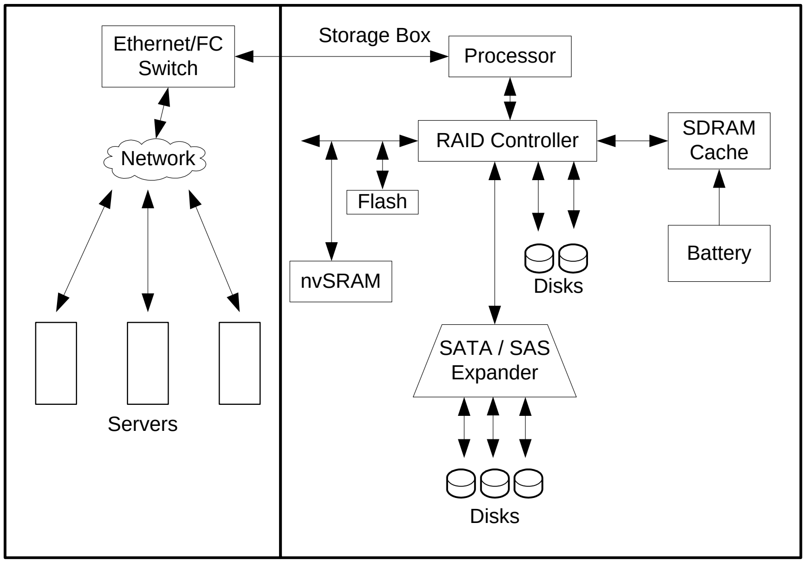

Figure 5. RAID Storage Architecture

The RAID controller has direct access to the cache memory, which enables fast read/write access to the storage system. The cache writes data in transition. The RAID system uses a cache to speed up the apparent I/O performance of the storage system so the host processor is not busy. In modern RAID systems, these cache memories are very large and are implemented with SDRAMs.

The write access from the host is acknowledged as soon as the data is received and is stored in the cache. Because disk drive access is not required, the throughput of a write operation increases significantly. The RAID controller transfers data from the cache to the disk drives. This operation, which takes place in the background, is transparent to the host, which assumes that the write operation is complete as soon as it receives an acknowledgment from the RAID controller. Although this approach significantly increases write performance, the inherent transparency raises a potential risk to data integrity if a power failure occurs after the host is acknowledged but before the data is committed to the disk. A write journal–based system is effective at recovering data from power losses.

The data stored is usually secured using a battery-backed cache that retains the contents of the cache even during power failure. On the next power-up, the storage system has to backtrack through transactions to determine which blocks of data were not written into the physical drives. Therefore, even with battery-backed cache, the system recovery time may be high, because the system needs to be checked for the exact failing point before recovery can begin.

To avoid recovery delays, many storage systems log data transactions from the host to the cache and from the cache to the disk in a nonvolatile circular buffer. When the power goes off, this log can be played back to precisely identify the state of the storage system before failure. This strategy, called “write journaling,” results in a write-in-progress record. nvSRAM is the preferred NV solution for write journaling because it is the fastest nonvolatile RAM. It also has proven reliability due to its SONOS architecture and has unlimited read/write endurance.

Smart Meters

Solid-state electricity meters primarily provide energy or power measurement and data display. They may also require battery operation to support features such as reading without power when the product is not connected to a power main. As the implementation and architecture become more sophisticated, electricity meters require additional processing power, larger flash memories for software stacks to support communication protocols, and larger nonvolatile memories for data logging and remote firmware updates. Smart meters offer additional functionality to energy meters, including real-time or near-real-time reads, power outage notification, and power quality monitoring. They allow price setting agencies to introduce different prices for consumption based on the time of day and season. These price differences help bring down peaks in demand (load shifting or peak lopping), thus reducing the need for additional power plants. Another type of smart meter uses nonintrusive load monitoring to automatically determine the number and type of appliances in a household, the amount of energy each appliance uses, and the times of operation. Utility companies use smart meters to conduct surveys on energy use. Figure 6 shows a typical block diagram of an energy meter.

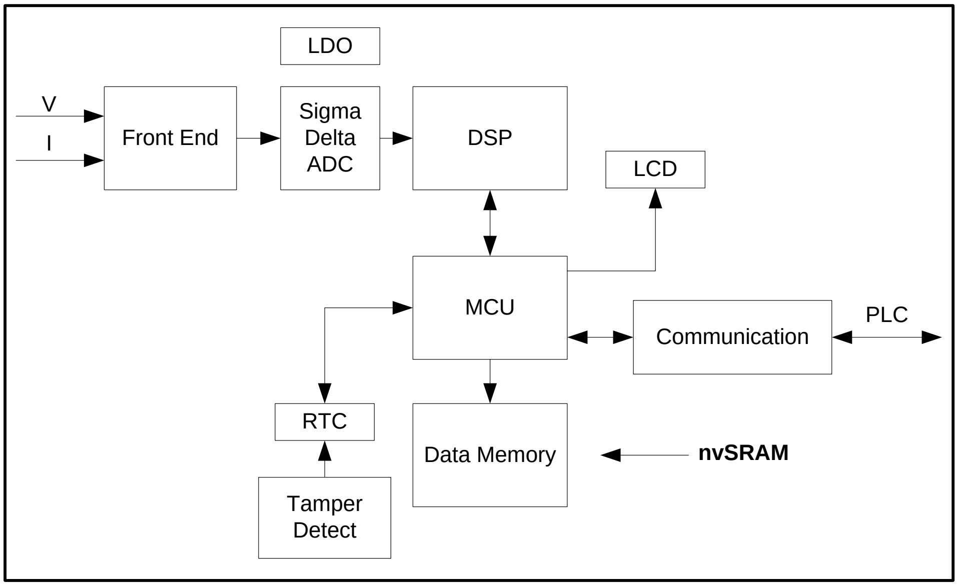

Figure 6. Advanced Energy Meter Block Diagram

The serial nvSRAM offers an industry-standard SPI interface with a speed of up to 104 MHz. Data write and read to the nvSRAM cells happen at the speed of the processor bus. This allows controllers to write into nvSRAM cells directly rather than storing first into an on-chip SRAM buffer and then transferring to external memory such as EEPROM or flash due to its slow page-based access. Also, the nvSRAMs eliminate the huge burden of wear leveling required by EEPROMs.

The serial nvSRAMs are available in the industry-standard 8-SOIC package. They have proven high reliability based on the standard SONOS process.

Summary

Cypress parallel nvSRAMs offer the functionality of fast asynchronous SRAMs with the feature of nonvolatility. Cypress nvSRAMs retain data using integral EEPROM providing high reliability. A small external capacitor (V_{CAP}) is used to transfer the contents of SRAM to the nonvolatile elements on power-down, which eliminates the need for using an external battery. The nvSRAMs have the same address, data, and control interface as asynchronous SRAMs, thereby simplifying use of the device. Other features are also added to the device internally to help applications reduce component count in their system. For example, some nvSRAMs are equipped with a real-time clock (RTC), thereby replacing components in many applications that require nvSRAM and RTC. In addition, Cypress serial nvSRAMs provide industry-standard I^2C and SPI interfaces, enabling low pin count applications with a high-speed serial interface. Cypress also offers asynchronous NAND interface nvSRAM which is aligned to a majority of the ONFI 1.0 specifications and supports data access speed up to 33 MHz. For more information on RTC, SPI, I^2C, and ONFI nvSRAM, refer to application notes AN61546, AN64574, AN74875, and AN91206.

Table 1 lists nvSRAM features and parameters. For more information, refer to specific datasheets for individual devices. See www.cypress.com/?id=65&tabID=72663 for the complete list of parts.

Table 1. nvSRAM Features and Parameters

Feature/Parameter		Specification (0.13 µm process technology)
V_{CC}	Parallel interface	5 V, 3 V, separate I/O (1.8 V)
	Serial interface	5 V, 3 V, 2.5 V
	NAND interface	3 V core, 1.8 V I/O
Device interface		Parallel, x8/x16
		Serial - I^2C, SPI
		Asynchronous NAND Interface (ONFI 1.0)
Speed	Parallel interface	20 ns access time
	Serial - I^2C	3.4 MHz
	Serial - SPI	104 MHz
	NAND interface	30 ns access time
Density	Parallel interface	Up to 16 Mb
	Serial interface	Up to 1 Mb
	NAND interface	16 Mb
Low voltage trigger level (V_{SWITCH})	3 V	2.65 V
	5 V	4.4 V
	2.5 V	2.35 V
Software functions: STORE, RECALL, AutoStore Enable/Disable	Parallel interface	Six address software sequences
	Serial interface	Opcodes
	NAND interface	Commands
RTC	Parallel interface	Available
	Serial interface	Available
	NAND interface	Not available
STORE cycle duration (t_{STORE})		8 ms
Power-up RECALL ($t_{HRECALL}$)	VCC = 3 V, 5 V	20 ms (30 ms for 16 Mb)
	VCC = 2.5 V	40 ms
Software RECALL (t_{RECALL})	Parallel interface	200 µs
	Serial interface	600 µs
Time allowed to complete SRAM Write Cycle (t_{DELAY})		25 ns
SLEEP	Parallel interface	Available in 16 Mb only
	Serial interface	Available
	NAND interface	Not available
Time to enter sleep mode after issuing SLEEP instruction		8 ms
Time for nvSRAM to wake up from sleep mode	V_{CC} = 3 V, 5 V	20 ms (30 ms in 16 Mb)
	V_{CC} = 2.5 V	40 ms
Endurance	SRAM read/write	Infinite
	NV STORE	1 million cycles
Data retention		20 years at 85 °C

Cypress Semiconductor
198 Champion Court
San Jose, CA 95134-1709
Phone: 408-943-2600
Fax: 408-943-4730
www.cypress.com

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CYPRESS MAKES NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS DOCUMENT OR ANY SOFTWARE OR ACCOMPANYING HARDWARE, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. To the extent permitted by applicable law, Cypress reserves the right to make changes to this document without further notice. Cypress does not assume any liability arising out of the application or use of any product or circuit described in this document. Any information provided in this document, including any sample design information or programming code, is provided only for reference purposes. It is the responsibility of the user of this document to properly design, program, and test the functionality and safety of any application made of this information and any resulting product. Cypress products are not designed, intended, or authorized for use as critical components in systems designed or intended for the operation of weapons, weapons systems, nuclear installations, life-support devices or systems, other medical devices or systems (including resuscitation equipment and surgical implants), pollution control or hazardous substances management, or other uses where the failure of the device or system could cause personal injury, death, or property damage (“Unintended Uses”). A critical component is any component of a device or system whose failure to perform can be reasonably expected to cause the failure of the device or system, or to affect its safety or effectiveness. Cypress is not liable, in whole or in part, and you shall and hereby do release Cypress from any claim, damage, or other liability arising from or related to all Unintended Uses of Cypress products. You shall indemnify and hold Cypress harmless from and against all claims, costs, damages, and other liabilities, including claims for personal injury or death, arising from or related to any Unintended Uses of Cypress products.