

Programmer's Guide for the Cypress HyperFlash Family

Author: Zhi Feng

Associated Part Family: HyperFlash

Associated Code Examples: None

Cypress HyperFlash memories use the low-signal-count, high-performance HyperBus™ interface. They are the industry's fastest NOR Flash devices that are able to transfer up to 333 Mbytes/s on the HyperBus interface. This document describes, for software programmers and system engineers, how to use the HyperFlash family.

Contents

1	Introduction.....	1	5	Reset.....	5
2	Basic HyperFlash Characteristics.....	1	6	CFI / Device ID.....	5
2.1	HyperFlash System Diagram.....	1	7	Maximizing Read Performance	6
2.2	Memory Architecture.....	2	8	Programming.....	7
2.3	Comparison to Traditional Parallel NOR Devices	2	9	Erasing	7
2.4	Comparison to SPI NOR Devices	3	10	Secure Silicon Region (SSR)	8
3	Configuring HyperFlash Devices	3	11	INT# Output Pin.....	8
3.1	Before First Memory Array Program or Erase.....	3	12	Conclusion.....	8
3.2	After All Configurations Are Set	4	13	References.....	8
4	Status Register.....	4		Document History.....	9
				Worldwide Sales and Design Support.....	10

1 Introduction


Cypress HyperFlash memories use the low-signal-count, high-performance HyperBus™ interface. They are the industry's fastest NOR Flash devices that are able to transfer up to 333 Mbytes/s on the HyperBus interface. This document describes, for software programmers and system engineers, how to use the HyperFlash family. For complete specifications of HyperFlash devices, see corresponding datasheets.

2 Basic HyperFlash Characteristics

The HyperFlash family consists of both 3.0-V KL-S and 1.8-V KS-S devices. Both 3-V and 1.8-V devices use the same software command set. HyperFlash combines the advantages of a low-signal-count interface memory such as SPI NOR with even higher read performance than Parallel NOR (PNOR) devices. This section compares HyperFlash to traditional PNOR and SPI NOR devices.

2.1 HyperFlash System Diagram

The following diagram shows how a typical embedded system chipset connects to a HyperFlash device. It requires the HyperBus memory controller that communicates with the HyperFlash device through a set of control and I/O pins. The software low-level driver provides basic read, write and erase functions to the upper-level applications.



2.2 Memory Architecture

HyperFlash has a uniform sector architecture with a sector size of 256 KB. A sector is the smallest erasable block in a flash device. To give users the flexibility of having smaller sectors, there is a user configuration option to partially overlay either the first or the last sector of the device. Those smaller sectors are called parameter sectors in which users typically store system parameters.

When programming to the memory, HyperFlash provides a Write Buffer, which is aligned on a 512-byte boundary. Therefore, the Write Buffer Programming command allows up to 512 bytes to be programmed in one operation.

Typically, device internal operations that take a period of time to complete, such as erasing a sector and programming, are called embedded operations (EO). During an EO, the device is busy and most commands are forbidden. Users can use the Read Status Register command to determine the completion of the EO.

Read commands are associated with pages. A page is a 16-word (32-byte) length and aligned unit internal to the device. Additional latency may be inserted when reading across the page boundaries. Refer to Section 7 for more details.

2.3 Comparison to Traditional Parallel NOR Devices

Software operations in HyperFlash are similar to traditional PNOR devices such as the Cypress GL-P and GL-S families. Although the physical pin connection of HyperFlash is different from PNOR devices, it adopts the PNOR command set as its baseline command set.

For example, to erase a sector of the device, users would issue a sequence of write commands identical to the erase command sequence used by the Cypress GL-P and GL-S families. Each command is written to the device through a different set of electrical signals and a signal protocol sequence compared to traditional PNOR. However, the electrical signaling difference is transparent to the software. The HyperFlash memory controller in the host system handles the translation of software write and read accesses into the HyperBus signaling protocol. Therefore, users can make use of the same Low Level Driver (LLD) software for earlier generations of PNOR with HyperFlash devices. HyperFlash devices have optional commands in addition to the baseline command set that enable additional features of the devices.

2.4 Comparison to SPI NOR Devices

The similarity between HyperFlash and SPI NOR devices is only in their signal pin physical placement on the package. The command format on the I/O bus for HyperFlash is completely different from SPI devices. As mentioned above, in terms of software command set, HyperFlash is backward-compatible with the Parallel NOR devices. That means the PNOR software low level driver can be reused on HyperFlash without changes.

3 Configuring HyperFlash Devices

HyperFlash devices provide a Nonvolatile Configuration Register (NVCR), and its counterpart, Volatile Configuration Register (VCR). The VCR is for users to temporarily change the configuration settings for testing. The VCR value will be reset to the one held in the NVCR at the next power cycle. If a nonvolatile configuration is desired, the NVCR should be updated to the desired value. The NVCR and VCR are collectively referred to as xVCR registers. [Table 1](#) shows the bit assignments of the xVCRs.

Table 1. Volatile and Non-Volatile Configuration Registers

xVCR Bit	Function	Settings (Binary)
xVCR.15	Reserved	1 - Reserved (default)
xVCR14 – xVCR12	Drive Strength	000 - (default) (Refer to datasheet for actual device dependent impedance)
xVCR.11	xVCR Freeze	0 - VCR or NVCR Locked (No Programming or Erasing of NVCR, no changes to VCR) 1 - VCR and NVCR Unlocked (factory default)
xVCR.10	SSR Freeze	0 - Secure Silicon Region Locked (Programming not allowed) 1 - Secure Silicon Region Unlocked (factory default)
xVCR.9 – xVCR.8	Parameter-Sector Mapping	00 - Parameter-Sectors and Read Password Sectors mapped into lowest addresses 01 - Parameter-Sectors and Read Password Sectors mapped into highest addresses 10 - Uniform Sectors with Read Password Sector mapped into lowest addresses. (factory default) 11 - Uniform Sectors with Read Password Sector mapped into highest addresses
xVCR.7 – xVCR.4	Read Latency	1011 - 16 Clock Latency (factory default) (Refer to datasheet for actual device dependent impedance)
xVCR.3	Reserved	1 - Reserved (default)
xVCR.2	Reserved	0 - Reserved (default)
xVCR.1 – xVCR.0	Burst Length	00 – Reserved 01 - 64 bytes 10 - 16 bytes 11 - 32 bytes (factory default)

3.1 Before First Memory Array Program or Erase

The HyperFlash family has default uniform sector architecture with 256 KB sector size. A user configuration option is available to partially overlay either the first sector or the last sector with eight 4-KB parameter sectors. This configuration is controlled by the Non-Volatile Configuration Register Bit 9-8. See the above table for the actual setting values.

If users want to configure parameter sectors, these two bits should be programmed to the desired values before any programming or erasing of the Flash array; otherwise, the data in the overlay portion of the Flash array may be lost.

3.2 After All Configurations Are Set

After the power-up default configuration has been determined and fully debugged, users may use the NVCR Bit 11 to permanently lock the xVCR registers. After the bit is programmed to 0, no changes can be made to xVCRs. The device configuration is permanently locked.

4 Status Register

HyperFlash has a single 16-bit Status Register (SR) that can be used to check the current state of the device, embedded operation status, or previous erase status. Only the lower eight bits of the SR are defined.

Table 2. Status Register


Bit #	15:9	8	7	6	5	4	3	2	1	0
Bit Description	Reserved	Reserved	Device Ready Bit	Erase Suspend Status Bit	Erase Status Bit	Program Status Bit	Write Buffer Abort Status Bit	Program Suspend Status Bit	Sector Lock Status Bit	Sector Erase Status Bit
Bit Name			DRB	ESSB	ESB	PSB	WBASB	PSSB	SLSB	ESTAT
Reset Status	X	0	1	0	0	0	0	0	0	0
Busy Status	Invalid	Invalid	0	Invalid	Invalid	Invalid	Invalid	Invalid	Invalid	Invalid
Ready Status	X	X	1	0 = No Erase in Suspension 1 = Erase in Suspension	0 = Erase Successful 1 = Erase Fail	0 = Program Successful 1 = Program Fail	0 = Program Not Aborted 1 = Program Aborted during Write to Buffer Command	0 = No Program in Suspension 1 = Program in Suspension	0 = Sector Not Locked during Operation 1 = Sector Locked Error	0 = Sector Erase Status Command Result = previous erase did not complete successfully 1 = Sector Erase Status Command Result = previous erase completed successfully

After issuing a command that triggers an embedded operation, such as programming or erasing a sector, users should always check the SR to make sure that the embedded operation has completed before proceeding to the next command. During an embedded operation, only the program/erase suspend¹ command or the Status Register Read command will be accepted. All other commands are ignored.

If an error happens during an embedded operation, users need to clear the error bit by using the Clear Status Register command before proceeding to the next command.

The following diagram shows a polling function that uses the Status Register to determine the device status after starting an embedded operation such as programming or erasing a sector. The algorithm here has the assumption that the type of the most recent operation, shown in the four large diamonds in the middle of the diagram, is passed into this software polling function. This implementation simplifies the software flow. It is feasible to design an algorithm that does not need to know the type of operation it is polling for; however, it will be more complicated.

¹ During program or erase operation, users may issue a suspend command to suspend the EO in order to quickly go back to read mode. The operation will be suspended till the resume command is entered.



5 Reset

There are three kinds of resets in HyperFlash device: Power on Reset (POR), also called cold reset; Hardware Reset, also called warm reset triggered by the RESET# signal; and the Software Reset triggered by the Software Reset Command (F0h).

Upon cold and warm resets, all VCR bits are loaded from the corresponding NVCR bit values.

The RSTO# pin can be used on the system to indicate the completion of a POR. The pin will transition from low to high state after the device completes the power-up stage, plus a user defined timeout period.

Software Reset is typically used to clear the Status Register or get the device back to Read Array Mode from an ASO state or an unknown state. Software reset will not affect ongoing embedded operations, or any Configuration Register values.

6 CFI / Device ID

Similar to earlier PNOR devices, the HyperFlash family follows the JEDEC Common Flash Interface (CFI) specification to provide a standardized data structure that may be read by a command. The data structure contains information such as various electrical and timing parameters, and special functions supported by the device. Software support can then be device-independent, Device ID-independent, and forward-and-backward-compatible for entire Flash device families.

Both ID (Autoselect) and CFI commands in HyperFlash allow users to access the combined ID/CFI data set. All data contained within the ID/CFI data set is available after using either the ID (Autoselect) or CFI Entry sequence.

See the HyperFlash family datasheet (referenced at the end of the App Note) for a complete explanation of the ID/CFI data structure.

7 Maximizing Read Performance

The HyperFlash memory array initial access time is defined as t_{ACC} at about 96 ns. This time is required to move the data from the Flash array to the data I/O; t_{ACC} is independent of the clock frequency the user has chosen to run at. The higher the clock frequency, the more clock cycles are required for the data to reach the device outputs. This time duration in terms of clock cycles is called latency clocks.

In HyperFlash, the number of latency clocks is controlled by four Latency Code Configuration Register bits: xVCR Bit 7-4. The following table shows the relationship between the latency code, latency clocks, and the maximum clock frequency that will satisfy the t_{ACC} requirement for initial read latency.

Table 3. Latency Settings

Latency Code	Latency Clocks	Maximum Operating Frequency (MHz)
0000	5	52
0001	6	62
0010	7	72
0011	8	83
0100	9	93
0101	10	104
0110	11	114
0111	12	125
1000	13	135
1001	14	145
1010	15	156
1011	16	166
1100	Reserved	NA
1101	Reserved	NA
1110	Reserved	NA
1111	Reserved	NA

Notes:

1. Default NVCR latency setting when the device is shipped from the factory is 16 clocks.
2. The Latency Code is the value loaded into (Non) Volatile Configuration Register bits xVCR[7:4].
3. Maximum Operating Frequency assumed to be using a device with $t_{ACC} = 96$ ns.

In order to maximize the read performance, users would want to read the maximum length of contiguous data in a single read command sequence to avoid multiple initial latency periods. That means the host controller is preferable to hold CS# LOW until all data is read. In order to achieve this, users cannot simply use `memcpy()` or similar read functions that reads one word at a time. Users will need to implement a controller-specific function to read multiple words of contiguous data for the application-level software. This function should accept data length as one of the parameters and have the host memory controller-specific register settings necessary for the memory controller to send the correct signal protocol sequence to the HyperFlash memory.

In HyperFlash, one full read page is 32 bytes. Users can configure the burst length of a read operation to be 16-, 32-, 64-byte wrapped sequence, or in a linear read sequence.

If the read is configured to be linear and the read address does not start on a 16-byte alignment boundary (0h or 8h word address multiple offset), in addition to the initial latency, the user may need to insert a page-crossing latency when crossing the first 32-byte alignment boundary (to the second page). The initial and first page-crossing latency cycles depend on the clock frequency and the starting address offset. See tables in the datasheet for the actual counts.

Note that all subsequent page crossings within the same linear access will not require any additional latency clocks. Wrapped read transactions of 16-byte and 32-byte bursts do not cross page boundaries and do not incur inter-page-boundary-crossing latencies. For a 64-byte wrapped burst read, additional latency may need to be inserted during the page boundary crossing, depending on the starting address.

8 Programming

HyperFlash family devices have a 512-byte write buffer, which is aligned on 512-byte boundaries. Programming data to the Flash is most efficient when writing in 512-byte length and aligned increments. Although smaller writes are allowed, for best performance, software should, whenever possible, program data in full, address-aligned, write buffer increments.

In HyperFlash, the Flash memory array is structured in 16-byte length and aligned Half-Pages. While multiple program operations within a Half-Page are not recommended, they are allowed for backwards compatibility with traditional Parallel NOR products. Programming more than once within the same Half-Page, per erase of the sector in which the Half-Page resides, reduces the data integrity of the Half-Page. If it is required to program in less than Half-Page multiples such that Half-Pages are programmed more than once per erase, it is recommended to add software error correction information for the data in such Half-Pages.

For example, a simple Flash file system may write 512-byte file records, each with 30-byte metadata. If the user programs the next file sector immediately after the 30-byte metadata, it would cause misalignments to all subsequent sectors, as shown in Table 4. In this case, it would be highly recommended that the user insert 2-byte padding data at the end of the 30-byte metadata structure so that all files sectors and metadata are aligned with the internal 16-byte Half-Pages, and 512-byte Write Buffer pages.

Table 4. Misaligned Data Storage

File sector	1st record 512 Bytes			1st metadata 30 Bytes		2nd record 512 Bytes				2nd metadata #0 Bytes	
Flash Page	HP 0	HP 31	HP 32	HP 33	HP 34	HP 64	HP 65	HP 66	

Note: HP: Half Page (16-byte)

Table 5. Aligned Data Storage

File sector	1st sector 512 Bytes			1st metadata 30 + 2 pad Bytes		2nd sector 512 Bytes				2nd metadata 30 + 2 pad Bytes	
Flash Page	HP 0	HP 31	HP 32	HP 33	HP 34	HP 65	HP 66	HP 67	

9 Erasing

To erase a sector in a HyperFlash device, issue the Erase command stated in the datasheet with the sector base address.

HyperFlash family devices have uniform sector size 256 KB. The first (Sector 0) or the last sector can be partially overlaid by eight parameter sectors each with 4-KB size. Users need to pay attention to the fact that in such overlay configuration, there is a sector above or below the parameter sectors (224 KB in size) between the parameter sectors and the remaining uniform 256-KB sectors. When issuing an erase command, the address used must be within the intended sector.

10 Secure Silicon Region (SSR)

HyperFlash family devices provide 1024 bytes of Secure Silicon Region (SSR) that is a One Time Programmable (OTP) area separated from the main Flash array. The area is divided into 32 individually lockable, 32-byte aligned regions. (32×32 bytes = 1024 bytes). Users can access the SSR by entering the SSR Address Space Overlay (ASO).

When reading from the SSR, enter SSR ASO and perform a read to an SSR address offset. If the address entered is outside of the 1024-byte SSR address range, main array data will be retrieved.

When programming the SSR, enter SSR ASO and perform Write Buffer or Word programming the same as in normal array programming.

The SSR is protected by the FREEZE bit in xVCR Bit10. If FREEZE is set, the program command will be ignored. No error is reported.

The Region 0 of the SSR (first 32 bytes) is a special region. The first 16 bytes of Region 0 are reserved for Cypress to program in a Random Number that can be used as a unique device identification such as a serial number. The next four bytes are the Lock Bits. Each lock bit controls the corresponding 32 SSR regions, from Region 0 to Region 31. Users can program these bits to individually lock any SSR regions. Once the region is locked, it is permanently locked.

Any attempt to program to the Random Number area will result in a program error. If an SSR region is locked by its lock bit, any attempt to program into the region will result in a program error.

When programming the SSR, the program page size is the same as the normal Flash array page program; that is, 512 bytes. That means the user can program multiple SSR regions in the same program command.

If the program data entered is more than the page size, the data will be wrapped to the beginning of the page, just like a normal page programming command. In this case, the wrapped data may coincide with Region 0, which is a special region as mentioned above. In this case, the program command may fail if it contains some data intended for the first 16 bytes of Region 0. Loading data beyond the end of the page programming buffer is not recommended.

11 INT# Output Pin

HyperFlash family devices offer an INT# output pin that users can configure to generate an interrupt to the system when transitioning from BUSY to READY state, for example, completion of an embedded operation.

Users can read the Interrupt Status Register (ISR) to verify a Busy-to-Ready event has occurred, or that the last POR was completed successfully, and then reset the ISR by writing the corresponding bits to 1.

12 Conclusion

Cypress HyperFlash family offers the same physical footprint as SPI devices, and the same software command interface as the Parallel NOR devices. It also combines advantages of low pin count of the SPI devices and a high throughput of NOR devices. In fact, its performance is even higher than Parallel NOR devices.

The HyperFlash family provides easy transition for users who have used SPI NOR or Parallel NOR devices in their systems, reduces the system cost, and helps achieve better system performance.

Contact Cypress Customer Support for additional help when using HyperFlash family devices.

13 References

- [S26KL/KSxxxS datasheet](#)
- [S26KL/KSxxxS Low Level Driver](#)

Document History

Document Title: AN99195 - Programmer's Guide for the Cypress HyperFlash Family

Document Number: 001-99195

Revision	ECN	Orig. of Change	Submission Date	Description of Change
**	4852549	ZHFE	08/07/2015	New Application Note

Worldwide Sales and Design Support

Cypress maintains a worldwide network of offices, solution centers, manufacturer's representatives, and distributors. To find the office closest to you, visit us at [Cypress Locations](#).

Products

Automotive	cypress.com/go/automotive
Clocks & Buffers	cypress.com/go/clocks
Interface	cypress.com/go/interface
Lighting & Power Control	cypress.com/go/powerpsoc
Memory	cypress.com/go/memory
PSoC	cypress.com/go/psoc
Touch Sensing	cypress.com/go/touch
USB Controllers	cypress.com/go/usb
Wireless/RF	cypress.com/go/wireless

PSoC® Solutions

psoc.cypress.com/solutions
[PSoC 1](#) | [PSoC 3](#) | [PSoC 4](#) | [PSoC 5LP](#)

Cypress Developer Community

[Community](#) | [Forums](#) | [Blogs](#) | [Video](#) | [Training](#)

Technical Support

cypress.com/go/support

PSoC is a registered trademark and PSoC Creator is a trademark of Cypress Semiconductor Corp. All other trademarks or registered trademarks referenced herein are the property of their respective owners.



Cypress Semiconductor
 198 Champion Court
 San Jose, CA 95134-1709

Phone : 408-943-2600
 Fax : 408-943-4730
 Website : www.cypress.com

© Cypress Semiconductor Corporation, 2015. The information contained herein is subject to change without notice. Cypress Semiconductor Corporation assumes no responsibility for the use of any circuitry other than circuitry embodied in a Cypress product. Nor does it convey or imply any license under patent or other rights. Cypress products are not warranted nor intended to be used for medical, life support, life saving, critical control or safety applications, unless pursuant to an express written agreement with Cypress. Furthermore, Cypress does not authorize its products for use as critical components in life-support systems where a malfunction or failure may reasonably be expected to result in significant injury to the user. The inclusion of Cypress products in life-support systems application implies that the manufacturer assumes all risk of such use and in doing so indemnifies Cypress against all charges.

This Source Code (software and/or firmware) is owned by Cypress Semiconductor Corporation (Cypress) and is protected by and subject to worldwide patent protection (United States and foreign), United States copyright laws and international treaty provisions. Cypress hereby grants to licensee a personal, non-exclusive, non-transferable license to copy, use, modify, create derivative works of, and compile the Cypress Source Code and derivative works for the sole purpose of creating custom software and or firmware in support of licensee product to be used only in conjunction with a Cypress integrated circuit as specified in the applicable agreement. Any reproduction, modification, translation, compilation, or representation of this Source Code except as specified above is prohibited without the express written permission of Cypress.

Disclaimer: CYPRESS MAKES NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Cypress reserves the right to make changes without further notice to the materials described herein. Cypress does not assume any liability arising out of the application or use of any product or circuit described herein. Cypress does not authorize its products for use as critical components in life-support systems where a malfunction or failure may reasonably be expected to result in significant injury to the user. The inclusion of Cypress' product in a life-support systems application implies that the manufacturer assumes all risk of such use and in doing so indemnifies Cypress against all charges. Use may be limited by and subject to the applicable Cypress software license agreement.