



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.



THIS SPEC IS OBSOLETE

Spec No: 002-03230

Spec Title: AN203230 - CYPRESS S25FL-S MULTI-I/O DDR
DLP OPTIMIZES READ PERFORMANCE AND
RELIABILITY

Replaced by: 002-29091

Cypress S25FL-S Multi-I/O DDR DLP Optimizes Read Performance and Reliability

This application note highlights the SPI flash DDR Quad I/O access protocol and how the use of DLP can optimize performance and reliability.

1 Abstract

Today's embedded systems are typically have larger code and data densities, faster start-ups, higher application performance requirements, while at the same time trying to reduce overall system cost. Cypress understands these conflicting constraints and continues its product innovations to offer best in class NVM solutions to address the next generation system requirements. For the past decade Cypress innovations have broaden our portfolio focus beyond parallel NOR's Asynchronous, Page, Burst interfaces to lower pin count, high read performance, low cost NOR SPI based solutions.

SPI interfaces have progressed from a single-bit, SDR, unidirectional input and output (x1) interface to SDR/DDR four-bit, bidirectional (x4) interface. Today's several leading NOR-based SPI memories achieved SDR 133-MHz clock rates and DDR quad I/O interface to facilitate a 66 MB/s continuous read throughput utilizing legacy SPI timing modes. With the addition of DDR timing approaches and the application of a DLP (data learning pattern), the Cypress S25FL-S SPI Family improves read throughputs by an additional 20 percent.

This application note highlights the SPI flash DDR Quad I/O access protocol and how the use of DLP can optimize performance and reliability.

2 Multi IO SPI Data Learning Pattern

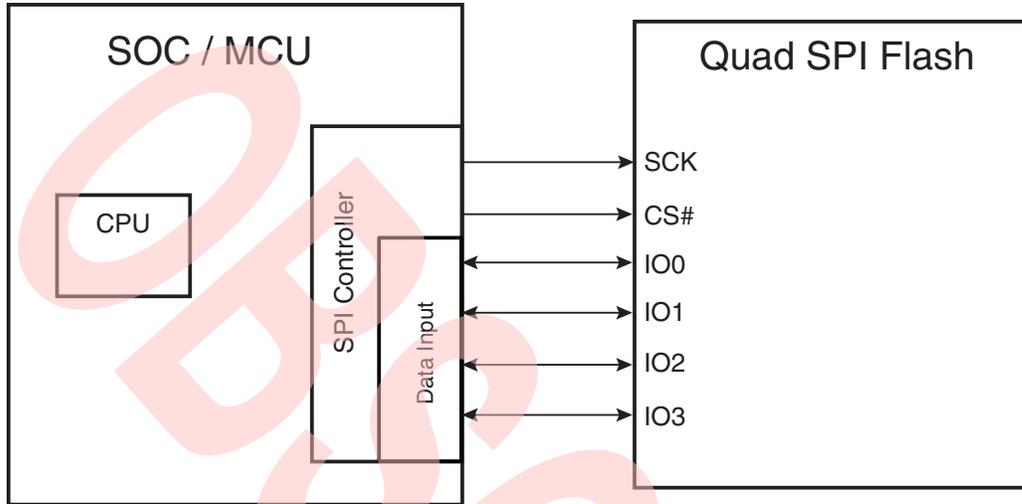
Today high speed embedded systems have more complex OS and Application requirements which typically results in increased read bandwidth requirements to provide acceptable performance at a neutral or lower cost point. Multi I/O SPI-based flash is becoming a pseudo- industry standard and its key features are its low pin count serial interface and bandwidths that are comparable to today's higher pin count parallel flash devices. The Cypress NOR SPI-based S25FL-S Eclipse™ family offers Multi I/O read access capability; offering excellent tradeoffs in reducing interface pin count and best in class read bandwidth. The S25FL-S optimized the base Multi I/O SPI-based read performance with the addition of DDR timing and the application of a data learning pattern (DLP). The S25FL-S improve read throughputs by an additional 20 percent.

The following sections highlights the SDR Multi I/O Quad read mode operation and a provides a comparison against DDR Multi I/O Quad DLP.

2.1 Quad Read Overview SDR

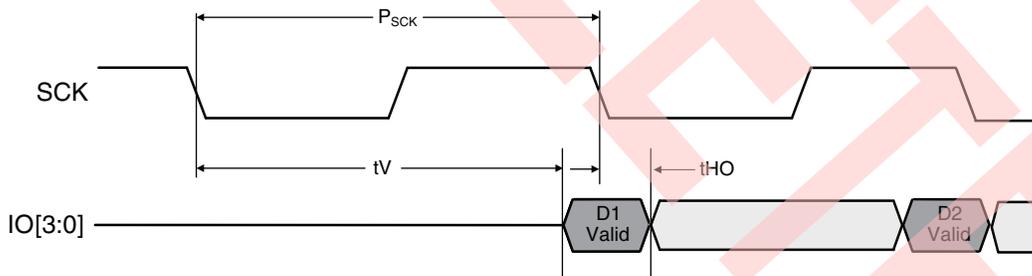
The Quad I/O SPI interface retains backward compatibility to support legacy x1 and x2 peripheral products. See Figure 1.

Figure 1. High Level Quad I/O Master / Slave Interface



The SDR (single-edge data rate) Multi I/O SPI read timing mode outputs a new data value upon each falling edge of SCK. After a period known as the clock-to-data-out time (t_V), data becomes valid and remains valid until shortly after the next falling SCK edge see Figure 2. The host typically uses this falling SCK edge to capture the data being output by the SPI flash. The hold time (t_{HO}) defines the length of time that data remains valid after a falling SCK edge.

Figure 2. Legacy SPI Timing $\geq t_{DV} = P_{sck} - t_V + t_{HO}$



Evaluating Data-Valid Time

Legacy SPI-based timing values, the size of the data valid (t_{DV}) is the clock period (P_{sck}) minus the time until data becomes valid (t_V) plus the hold time (t_{HO}) after the next falling clock edge: $t_{DV} = P_{sck} - t_V + t_{HO}$

- Example Legacy SDR Data Valid Window:
 - Clock Period: $P_{SCK} = 7.5 \text{ ns}$ (80 MHz)
 - Open: $t_V = 6.5 \text{ ns}$ / Close: $t_{HO} = 0 \text{ ns}$
 - Data Valid = $P_{SCK} - t_V + t_{HO}$: $7.5 \text{ ns} - 6.5 \text{ ns} + 0 \text{ ns} \sim 1 \text{ ns}$

If one has the assumption that t_V and t_{HO} timings are fixed then the data valid window compresses as the SCK frequency increases which limits SCK to $\sim 133 \text{ MHz}$.

In a SPI device the t_V and t_{HO} timing track each other; a device with a longer t_V have a longer t_{HO} and a device with a short t_V have a short t_{HO} . The data valid window timing varies with respect to the next falling clock edge and

t_V . Utilizing this information that t_V and t_{HO} track one another one can define the size of the data valid (t_{DV}) as equal to:

$$t_{DV} = P_{sck} - t_{O_SKEW} - t_{OTT}$$

The key benefit of understanding the t_V / t_{HO} timing constraint allows for optimization of flash data-valid window. Consider the 3V flash using a 133-MHz clock ($P_{sck}=7.5$ ns), a t_{O_SKEW} of 600 ps. and an Output slew rate of 2V/ns. The output rise/fall time (t_{OTT}) is:

$$t_{OTT} = V_{output_swing}/Output_slew_rate [4]= 3 V/(2 V/ns) = 1.5 ns$$

$$t_{DV} = P_{sck} - t_{O_SKEW} - t_{OTT} \\ = 7.5 ns - 600 ps - 1.5 ns = 5.4 ns$$

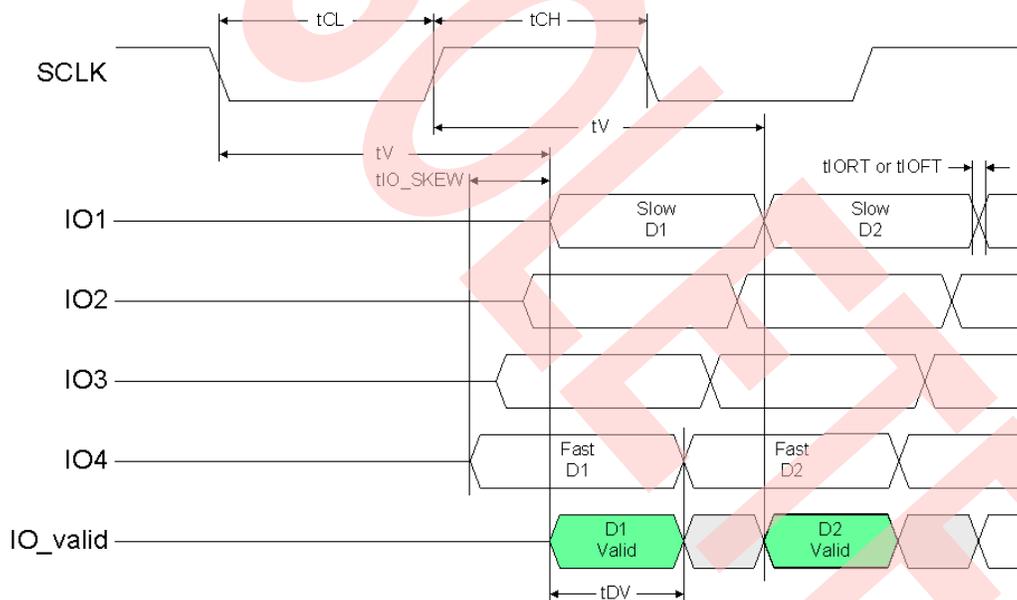
The new data-valid window can provide a significant improvement over the legacy SPI timing.

2.2 Multi I/O Quad Read DDR

These insights concerning the data valid window can be applied SPI DDR mode. In DDR mode the t_{O_SKEW} and t_{OTT} values do not change but new data is output every half clock cycle rather than after every full clock cycle, as is the case for SDR mode (see Figure 3). The t_{DV} is defined as:

$$t_{DV} = t_{CLH} - t_{O_SKEW} - t_{OTT} [5]$$

Figure 3. DDR Quad SPI Timing



This DDR SPI data-valid period for a system running at a given clock speed is identical to the data-valid period for an SDR system running at twice that clock speed. This means a DDR SPI device can reliably achieve the same read data rate at a significantly slower clock speed. For example, a QUAD DDR SPI with a clock operating at 80 MHz can achieve 80 MB/s. Note operating at a slower clock speed increases the data valid time.

Consider the 3V flash DDR Read access using a 80 MHz clock ($P_{sck}/2 = 6.25$ ns), a t_{O_SKEW} of 600 ps, and an Output slew rate of 2V/ns. The output rise/fall time (t_{OTT}) is:

$$t_{OTT} = V_{output_swing}/Output_slew_rate [4]= 3V/(2V/ns) = 1.5 ns. \text{ Assume 50\% SCK duty cycle.}$$

$$t_{DV} = P_{sck}/2 - t_{O_SKEW} - t_{OTT} \\ = 6.25 ns - 600 ps - 1.5 ns \sim 4.15 ns$$

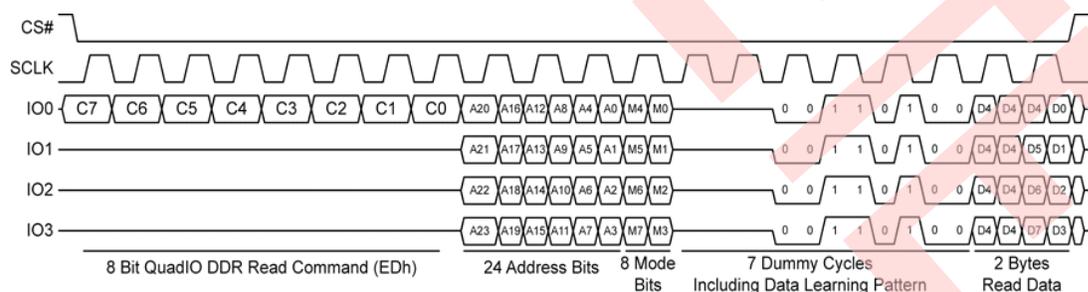
SPI-DDR Read Operation — Data Learning Patterns

In a DDR implementation, the t_V time can be greater than a half clock period which means a specific SCK edge cannot be used by the Master to reliably capture the data coming from the flash. The Master must skew the data-capture point with respect to each SCK edge in order to reliably capture data. Prior to detailing how the Master might facilitate appropriate sampling skews let's discuss the S25FL-S six new SPI DDR read operations. The new DDR read protocol is available for x1, x2, and x4 interfaces with either three- or four-byte addressing. Consider an SPI-DDR read operation performed using a Quad I/O interface with three-byte addressing (command EDh). The command sequence is much like a standard Quad I/O read operation with the exception that the address, mode, and data bits are transferred on both rising and falling clock edges (DDR) rather than the standard Quad I/O SDR protocol. The SPI-DDR read protocol is processed as follows:

1. The instruction (command operation code) is transferred in an SDR manner for compatibility with all other legacy SPI instructions. After the instruction is sent, all remaining transfers are DDR.
2. Target address is transferred (DDR).
3. Mode bits are loaded (DDR).
4. Read latency (dummy) cycles are issued while target data is extracted from the array.
5. A data learning pattern (DLP) is output by the SPI NOR flash during the last four dummy cycles (DDR).
6. Target data is output by the SPI NOR device (DDR).

The new SPI-DDR protocol as shown in Figure 4 adds an 8-bit data-learning pattern (DLP) that is output by the SPI flash using DDR protocol on the four dummy clock cycles just before output of the target data. The DLP provides a known data sequence on each data signal so that the host controller can determine the optimal capture timing to use when receiving the read data. The dummy cycles that carry the DLP occur during the idle period in the legacy Quad I/O read protocol while the target data is retrieved from the memory array; it can be accomplished without impacting performance. The DLP presents the same timing, phase delay, and skew characteristics that exist during output of the target data. The clock to data output delay (t_V in Figure 3) will be the same for the DLP and the target data on each individual data signal. Data phase delay and skew arise from issues related to either the memory device or the system environment. Memory device timing variations are caused by process, voltage, temperature, and output-to-output skew. System-level variations are introduced by PCB parasitics, trace-length mismatches, and bus capacitive loading. Collectively, these timing phase and skew relationships are represented in Figure 3. The timing characteristics of the known data learning pattern will allow the host controller to compensate for both the device- and system-level timing phase and skew offsets when valid data is present on the bus.

Figure 4. Quad IO SPI-DDR Showing DLP Read Transaction



Key Points:

- Data Learning Pattern output by memory
- Oversampled by host
- Optimal data capture point determined
- Data read from device
 - Calibration upon every read transaction
 - Provides Compensation for Process, Voltage, Temperature
- 160 Mbps (per IO) Data Rate Today
 - Strategy extendable to higher rates

2.3 Data Learning Pattern Storage and Definition

The non-volatile data learning register (NVDLR) and the volatile data learning register (VDLR) are used to define the sequence of DLP values (8 bits on each of the four Lows) that are used during an SPI-DDR read operation to train the host controller (see Figure 5). The NVDLR can be programmed one time (OTP) with a customer-specific DLP value. During power-up or reset, the value in the NVDLR is loaded into the VDLR. The sequence of values used as the DLP is defined in the VDLR during SPI-DDR read operations. The VDLR can be read and written directly by the host system. When the VDLR is 00h, the DLP will not be output during DDR read operations, thus providing an option to turn off the DLP.

The choice of an appropriate DLP is up to the system developer but the pattern should be chosen to maximize skews dependent on the bit-stream sequence. The most significant pattern-dependent skewing is bounded by transitions from states that have been stable for extended periods and states that have existed for shorter periods of time; for example, at higher frequencies, the High to Low transition of a signal behaves slightly different when the bit stream is xx110 than when the bit stream is xx010. The High reached in the xx110 pattern is usually a higher voltage than the High reached with the xx010 pattern. The higher starting voltage will mean that it will take slightly longer to reach a valid LOW state during a High to Low transition. The xx110 High to Low transition demonstrates a 'strong 1,' while the xx010 transition demonstrates a 'weak 1.'

The DLP pattern should be chosen to include at least one instance of:

- weak 0, strong 0, weak 1, and strong 1

One pattern fulfilling these requirements is 34h (00110100b). The edges in the 34h pattern step through the following transitions: Strong 0 -> Strong 1 -> Weak 0 -> Weak 1

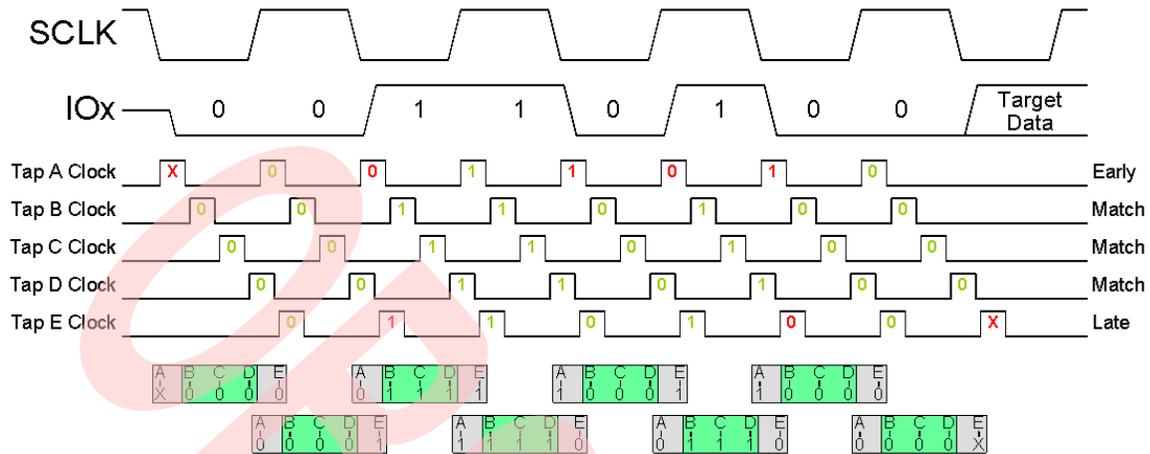
Any data learning pattern that includes these four transition types should maximize pattern-dependent skew characteristics.

2.4 Host Capture Strategy

The overall data-capture strategy for the host memory controller is to use the DLP input as a test sequence to characterize system response and determine tV and tDV. Once the data eye has been identified during the DLP portion of the DDR read sequence, the controller selects the optimal data-capture point to maximize the timing margin for the read data.

A common way to create the Master data-capture logic is via series of skewed data-capture points that span the data-valid window. The implementation for a single I/O might consist of five channels with a fixed sampling delay between each of the channels. The five delayed strobes (A through E) could be generated with a delay-locked loop (DLL) or using an oversampling clock that is in turn generated using an internally available higher frequency clock. The host controller samples the target I/O while the DLP is being output. The phase-delayed strobes (A-E) are triggered by the eight SCLK edges when the DLP is output.

Figure 5. Host Capture Strategy



Key Points:

- Data oversampled
- Samples from 'taps' B, C, and D always successfully capture the DLP
- In this case Tap C provides the greatest margin
- Use Tap C to capture data for the remainder of this read transaction
- Recalibration can be performed prior to every read transaction, provides more robust / reliable operation across operating conditions

3 Conclusion

As embedded applications continue to demand higher performance the legacy SPI interface and protocols must continue to accommodate higher read speeds. The DLP approach enables SPI-DDR NOR flash to move beyond today 133 Mbps (per pin) data rate. This new feature provides the embedded designers another attractive nonvolatile memory solution to maximize read data throughput while minimizing pin count, PCB complexity, package size, and cost. These principles and feature improvements apply across many market applications whether it is industrial, automotive graphics, or consumer; DLP-enabled SPI-DDR NOR flash provides another enhanced solution to improve system design and performance, all for a reasonable cost.

4 References

[Cypress: S25FL256S Data Sheet](#)

Cypress Article: Data learning and SPI boost NOR flash performance by Cliff Zitlaw

Document History Page

Document Title: AN203230 - Cypress S25FL-S Multi-I/O DDR DLP Optimizes Read Performance and Reliability				
Document Number: 002-03230				
Rev.	ECN No.	Orig. of Change	Submission Date	Description of Change
**	–	–	05/20/2015	Initial version.
*A	5041698	MSWI	12/08/2015	Updated to Cypress template.
*B	5870112	AESATMP8	09/01/2017	Updated Cypress Logo and Copyright.
*C	6743847	BWHA	12/16/2019	Obsolete document.

OBsolete

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

ARM® Cortex® Microcontrollers	cypress.com/arm
Automotive	cypress.com/automotive
Clocks & Buffers	cypress.com/clocks
Interface	cypress.com/interface
Internet of Things	cypress.com/iot
Memory	cypress.com/memory
Microcontrollers	cypress.com/mcu
PSoC	cypress.com/psoc
Power Management ICs	cypress.com/pmic
Touch Sensing	cypress.com/touch
USB Controllers	cypress.com/usb
Wireless Connectivity	cypress.com/wireless

PSoC® Solutions

[PSoC 1](#) | [PSoC 3](#) | [PSoC 4](#) | [PSoC 5LP](#) | [PSoC 6](#)

Cypress Developer Community

[Forums](#) | [WICED IOT Forums](#) | [Projects](#) | [Video](#) | [Blogs](#) | [Training](#) | [Components](#)

Technical Support

cypress.com/support

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

© Cypress Semiconductor Corporation, 2015–2019. This document is the property of Cypress Semiconductor Corporation and its subsidiaries, including Spansion LLC (“Cypress”). This document, including any software or firmware included or referenced in this document (“Software”), is owned by Cypress under the intellectual property laws and treaties of the United States and other countries worldwide. Cypress reserves all rights under such laws and treaties and does not, except as specifically stated in this paragraph, grant any license under its patents, copyrights, trademarks, or other intellectual property rights. If the Software is not accompanied by a license agreement and you do not otherwise have a written agreement with Cypress governing the use of the Software, then Cypress hereby grants you a personal, non-exclusive, nontransferable license (without the right to sublicense) (1s) under its copyright rights in the Software (a) for Software provided in source code form, to modify and reproduce the Software solely for use with Cypress hardware products, only internally within your organization, and (b) to distribute the Software in binary code form externally to end users (either directly or indirectly through resellers and distributors), solely for use on Cypress hardware product units, and (2) under those claims of Cypress's patents that are infringed by the Software (as provided by Cypress, unmodified) to make, use, distribute, and import the Software solely for use with Cypress hardware products. Any other use, reproduction, modification, translation, or compilation of the Software is prohibited.

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.

Cypress, the Cypress logo, Spansion, the Spansion logo, and combinations thereof, WICED, PSOC, CapSense, EZ-USB, F-RAM, and Traveo are trademarks or registered trademarks of Cypress in the United States and other countries. For a more complete list of Cypress trademarks, visit cypress.com. Other names and brands may be claimed as property of their respective owners.